Packetshaper 11.9

Symantec
A Division of Broadcom

TOC

Overviews - The Big Picture 50

Access-Link Monitoring Overview 51

Access-Link Monitoring Requirements 52

MIB Variables Polled by the Access-Link Monitoring Feature 52

Adaptive Response Overview 53

Introduction 53

Adaptive Response Agents 53

Adaptive Response Dashboard 55

Drill-Down Incident Reports 56

Adaptive Response Action Files Overview 58

Including Automatic Notification in an Action File 58

Creating an Action File 59

Action File Variables 59

Variables Available to All Agents 59

Agent-Specific Variables 60

User-Defined Variables 64

Adaptive Response Agents Overview 66

BCAAA Overview 67

How Does PacketShaper Work with BCAAA? 67

User Caches 68

Planning for your BCAAA Installation 68

Bypass Mode Overview 69

Captive Portal Overview 70

Traffic Classification Overview 73

Terms 73

Why Classify Traffic? 73

Creating Traffic Classes 73

Control Strategy Overview 74

Other Useful Considerations 75

Customer Portal Overview 76

Overview of Event Notification 78
Configuring the Event Feature 78
Setting Up and Managing Events 78

Flow Detail Records Overview 79
What Type of Information is in a Flow Detail Record? 79
Flow Detail Record Collectors 80
How Often are Flow Detail Records Emitted? 80
End-of-flow FDRs *80*
Intermediate FDRs *81*
Classification of Flow Detail Records 81
Feature Requirements and Limitations 82

IntelligenceCenter Overview 83

IPv6 Classification Overview 84
Control of IPv6 Traffic 84
Features that Support IPv6 84
Feature Limitations 85

Licensing Overview 86

Matching Rules 87
What is a Matching Rule? 87
Device 88
URL Category 88
Service Group 89
Service 89
Operation 89
Protocol Family 89
Users 90
Server Location 90
Classifying by Server Location *90*
Example of Server Location *91*
Ports 91
Port Number Range *91*
Non-Contiguous Port Numbers *91*
Proxy this Service to a Non-Standard Port 92
Host/Subnet 92
Name *92*
IP Address *92*
Host List *93*
Subnet and Mask *93*

MAC Address *93*
Application-Specific Criteria **93**
Diffserv (DSCP) Matching Rule **94**
Diffserv Type: Code Point (DSCP) *94*
Diffserv Type: COS/TOS *94*
VLAN Identification **95**

Partition Overview **96**
Managing Bandwidth **96**
Protecting Traffic *96*
Limiting Traffic *96*
Dividing Capacity *97*
Assigning Bandwidth Dynamically *97*
Oversubscribing *97*
Initial Configuration **97**
Types of Partitions **97**
Accessing Available Bandwidth **98**
Hierarchical Partitions **98**
Hierarchical Partition Examples *99*

Policy Overview **100**
Policy Types **100**

PolicyCenter Overview **102**
What is the Benefit of PolicyCenter? *102*
PolicyCenter Units Must be in Shared Mode *103*

RTM Overview **104**
RTM Functions *104*
Benefits *104*
How Response Times are Tracked *104*

Sky UI Overview **106**
Advanced UI **107**

SNMP Overview **109**
Installation **109**
SNMPv1/v2c **109**
SNMPv3 ..... **109**
Supported SNMP MIBs **110**

Standby Overview **113**
Requirements **114**
Additional Information **115**

Traffic Tree Overview (Advanced UI) 116
Creating a Traffic Tree 116
Large Traffic Trees 117
Traffic Class Icons 118

Class Tree Overview (the Sky UI) 120
Understanding the Class Tree 120

URL Categories Overview 122
WebPulse Express 123
Category Classification 123
WebPulse Interaction 125
URL Cache 125
Feature Limitations 125

User Awareness Overview 127
Support for International Characters in User and Group Names 127
Requirements 127
Feature Limitations 128

Watch Mode Overview 129
Non-Inline Deployment Requirements and Limitations 129

Web Application Classification Overview 131
WebPulse Express 132
URL Cache 133
Legacy Built-In Web Applications 134
Feature Requirements 134
Feature Limitations 134

WebPulse Overview 135
WebPulse Express 135
URL Cache 136
Monitoring WebPulse-Related Traffic 137
WebPulse Deployment Requirements 138

Recommendations for Real-World Situations 139

Analysis and Monitoring Recommendations 140

Classify Network Traffic 141

Create an Application-Based Traffic Tree 143

Create a Category-Based Traffic Tree 145

Create a Category-Based Traffic Tree: to Find Out Which Categories Users are Accessing 146

Create a Category-Based Traffic Tree: to Manage Traffic from Specific Categories 147

Create a Category-Based Traffic Tree: to Treat Secure and Unsecure URL Categories Differently 150

Create a Category-Based Traffic Tree: to Manage Types of URL Categories 152

Use the Model Tree 153

Create a Simple Location-Based Traffic Tree 155

Create a Location-Based Traffic Tree with Per-Location Applications or Categories 156

Create a Location-Based Traffic Tree with Global Applications 159

Analyze your Link's Behavior 161

Assess Wasted Network Capacity 163

Identify Performance Saboteurs 166

Analyze an Application's Behavior 169

Analyze Bandwidth Usage Patterns 171

Evaluate an Application's Efficiency 173

Analyze Application Response Times 174

Monitor a Condition of Interest Notify Someone about a Condition of Interest Respond to a Condition of Interest ... All Automatically 177

Monitor and Respond to Conditions of Interest using Agent Templates 178

Phase One: Plan Your Adaptive Response Agent 178

Phase Two: Create your Action Files 180

Phase Three: Add or Edit your Agent 181

Monitor and Respond to My Own Custom Condition 183

Phase One: Plan Your Adaptive Response Agent 183

Phase Two: Create and Test your Action Files 185

Phase Three: Add Your Agent 186

Notify Someone of Situations of Interest 188

Define and Monitor a Set of Reports 189

Identify Mysterious Traffic 190

Identify New Traffic 193

Identify Users Who Access a Category 195

Sniff Without a Sniffer 196

Use Standardized, Artificial Traffic for Analysis 198

Control Strategy Recommendations 201

Protect Critical Application Performance 202

Block Unwanted Traffic 204

Block a URL Category 205

Contain a Greedy Business Application 206

Reserve Bandwidth 207

Limit an Application's Total Bandwidth 208

Subdivide Bandwidth 209

Insulate Users of the Same Application 210

Provision Bandwidth Equitably 211

Quarantine Bandwidth Abusers 214

Quarantine Bandwidth Abusers - Use Host Accounting 215

Control Branch Traffic from a Main Site 219

Creating Application Policies for Each Branch's Traffic to Inherit (for location-based traffic trees with global applications) *220*

Integrate ProxySG with PacketShaper 222

Recommendations for Per-Application Management Strategies 225

Manage Citrix Performance 226
Balance DICOM's Bandwidth Demands 228
Protect ERP Performance 230
Common Performance-Management Strategies for Healthcare Environments 231
Control Instant Messaging 233
Control Peer-to-Peer Downloads 235
Limit Employee Access to Recreational Web Applications 237
Control Users' Access to an Application 239
Control Internet Radio Traffic 241
Control Streaming Media 242
Manage Voice and Video Sessions 244
Restrict Skype Traffic 247
Prepare for the Next Napster 248
Security Recommendations 250
Security Best Practices for PacketShaper S-Series 251
Secure Interfaces *251*
Invisible IP Address *251*
Virtual, Non-Routable IP Address *251*
Local User Accounts *251*
Third-Party User Authentication *252*
Secure Unit Access Protocols *252*
Supported Ciphers and Algorithms for SSH *253*
Transport Layer Security for HTTPS *254*
Secure SNMP Monitoring *255*
Secure Configuration Backups *255*
Session Inactivity Timeouts *255*
Zeroization Before RMA *256*
Secure Management by PolicyCenter *256*
Secure Your WAN Applications 258
Application Flow Validation 258
Web Proxies *258*
Encryption *259*

Port Hopping *259*
Rogue Server Detection **259**
Top Talkers and Top Listeners *259*
Flow Detail Records *260*
Application Overload Protection **260**
Detection of SYN and Spoofing Attacks **261**
Detection of Recreational Traffic **261**
Protecting Against Network Congestion Events **261**
WAN Application Forensics **262**
Graphing Tools *262*
Flow Detail Records *262*
Host Analysis *263*
Packet Capture *263*

Detect and Limit DoS Attacks **265**
Monitor for Attacks **265**
Limit Flood Attacks (Back to top) **266**
Limit the Rate of New Traffic Flows to or from One Host *266*
Limit the Amount of ICMP Traffic *267*
Limit the Number of Concurrent Traffic Flows for One Traffic Class *268*
Detect Psuedo-Address Attacks (Back to top) **268**
Disrupt Virus Attacks (Back to top) **268**

Administration Recommendations **270**

Vary Management Strategy by Time and Day **271**

Migrate Legacy Configurations to PacketShaper S-Series **273**

Verify PacketShaper Health **275**
Available Capacity *275*
Diagnostic Commands *276*

Advanced UI Tasks - Step-by-Step **277**

Setup Tasks **278**

Run Guided Setup **279**

License PacketShaper **280**

Install Initial License **281**

Upgrade PacketShaper License **284**

Modify Settings 287

Reconfigure Basic Settings 289

Enable/Disable Traffic Shaping 293

Subscribe PacketShaper Units to PolicyCenter 294

Add a New PacketShaper to PolicyCenter 294

Add a Currently Operating PacketShaper to PolicyCenter 294

Confirm that your PacketShaper is Installed in PolicyCenter 295

Specify Security Settings 297

Password Recovery 301

Configure RADIUS 302

Configure RADIUS Servers 303

Configure RADIUS Logins to PolicyCenter 304

Configure RADIUS Authentication Service 305

Configure RADIUS Accounting Service 307

Log In and Out with RADIUS 309

Logging In with RADIUS 309

Logging Out 309

Configure TACACS+ 311

Configure Linux TACACS+ Servers using the Cisco TACACS+ Daemon 312

Configure the TACACS+ Server 312

Start the TACACS+ Server 313

Configure TACACS+ Authentication Service 314

Configure TACACS+ Accounting Service 316

Log In and Out with TACACS+ 317

Logging In with TACACS+ 317

Logging Out 317

Define Flow Detail Record Collectors 318

Enable/Disable an FDR Collector 320

Enable Intermediate FDRs ... 321
Remove an FDR Collector ... 322
Modify Date and Time Settings ... 323
Upgrade the PacketShaper Software Image ... 325
Recovery Feature ... 327
Configure an SMTP Server ... 329
Reset the Traffic Tree ... 330
Reboot the Unit ... 332
Reset Unit to Factory Default Settings ... 333
Configure SNMP on PacketShaper ... 334
Configure PacketShaper for SNMPv1-v2c Support ... 336
Configure PacketShaper for SNMPv3 Support ... 337
Configure SNMP Access Tables ... 337
Configure the Views Table ... 337
Configure the Access Groups Table ... 340
Configure the Users Table ... 343
Configure SNMP Notify Tables ... 346
Configure the Remote Users Table ... 346
Configure the Notify Targets Table ... 349
Adjust System Variables ... 354
Revert to Default System Variable Settings ... 360
Configure Heartbeat Emission ... 361
Configure URL Categorization ... 363
Manually Update the WebPulse Maps ... 364
Select Categories to be Discovered ... 366
Configure WebPulse ... 368
Configure User Awareness on PacketShaper ... 370

Install BCAAA 371
Planning for your BCAAA Installation 371
Download and Install BCAAA 372
Configure PacketShaper as a BCAAA Client 374
Verify BCAAA Server Configuration 376
Traffic Classification Tasks 378
Enable/Disable Traffic Discovery 379
Turn on Traffic Discovery for an Individual Class 380
Create a Traffic Class 382
Add a Matching Rule 384
Modify a Matching Rule 385
Delete a Matching Rule 386
Classify URL Categories 387
Create a Category-Based Class 388
Auto-Discover Category Classes 390
Discover Categories under Web-Based Classes 391
Create a Class for Multiple Categories 392
Classify Users 393
Create a User-Based Class 394
Create a User Group Class 396
Create a Class Based on a User List 398
Open the User List Editor 400
Create User Lists 401
Edit a User List 403

Delete a User List 405

Manually Create an IPv6 Class in the Advanced UI 406

Create a Class for an IPv6 Server 407

Create a Service-Centric Traffic Tree for IPv6 408

Create an IPv6-Based Branch in the Traffic Tree 409

Create a Site-Based Traffic Tree for IPv6 Subnets 410

Create a Citrix-ICA Class 411

Citrix Criteria 412

Classifying by Published Application 412

Classifying by Client Name 413

Classifying by Priority 413

Classify Web Traffic 414

Create Web Classes 415

Classify Web Applications 417

Manually Creating Web Application Classes 417

Autodiscovering Web Applications 418

View Lists of Web Applications and Operations 419

Create FTP Classes 421

Create a Database Class 422

Automatic Creation of Database Classes 422

Manual Creation of Database Classes 422

Classify Traffic Based on VLAN ID or Priority 424

VLAN Identification (802.1Q or ISL) 424

VLAN Priority (802.1p) 425

Classify Diffserv Code Point Traffic 426

Classify Diffserv COS/TOS Traffic 427

Classify ICMP Messages 429

Classify RTP Traffic 431
Service Group Tasks 433
Create a Custom Service Group 434
Move Services to a Different Group 436
Return Services to Default Groups 438
Restore Settings to Factory Defaults *438*
Move Services into their Default Groups *438*
Delete a Custom Service Group 440
Create a Service Group Class 441
Create a Service Group Class in the Advanced UI 442
Create a Service Group Class in the Sky UI 443
Traffic Tree Management Tasks 444
Collapse and Expand the Traffic Tree 445
Collapse the Tree 445
Expand the Tree 445
Set the Tree Depth 446
Large Traffic Trees 446
Create a Folder Class 447
Create an Exception Class 448
Create an Inheritable Traffic Class 449
Inherited Class Example 450
Update Class Properties 451
Test a Traffic Class 453
Copy a Class 455
Copy Multiple Classes 456
Delete a Class 457

Delete Multiple Classes ... 458
Move a Class ... 459
Monitoring Tasks ... 460
Monitor Traffic in the Advanced UI ... 461
Standby Feature ... *463*
Select and Sort ... 465
Selecting Columns to Display ... 465
Sorting the Monitor Traffic Table ... 466
Selecting Rows to Display ... 466
Reset Monitor Statistics ... 468
Update the Monitor Statistics ... 469
Track Hosts and Users that Generate the Most Traffic ... 470
List Top Hosts or Users ... 472
Displaying the List from the Monitor Tab ... 472
Displaying the List from the Manage Tab ... 473
Create a Class for a Top Host or User ... 475
Verify Traffic Classes ... 478
Partition Tasks ... 479
Create a Partition ... 480
Sizing a Static Partition ... 482
What is a Static Partition? ... 482
Choosing a Unit of Measurement ... 482
When to Use Explicit Partition Sizes ... *482*
When to Use Relative Partition Sizes ... *483*
Partition Examples ... 484
Partition for Music Downloads ... *484*
Partition for Citrix ... *484*
Partition(s) for Streaming Media ... *484*
Is the Partition Size Appropriate? ... 484
Modify a Partition ... 485

View Partition Information 486
List Classes with Partitions 486
View the Partition Summary 487
Create Partition Utilization Graphs 490
Delete a Partition 491
Create a Dynamic Partition 492
Modify a Dynamic Partition 496
Remove a Dynamic Partition 498
Enable/Disable Traffic Shaping 499
Policy Tasks 500
Use Suggested Policies 501
Set a Rate Policy in the Advanced UI 502
Set a Priority Policy in the Advanced UI 504
Priority Policy Examples 505
Set a Never-Admit Policy in the Advanced UI 506
Special-Case Never-Admit Considerations 507
Never-Admit Policy Example 509
Set an Ignore Policy in the Advanced UI 510
Set a Discard Policy in the Advanced UI 511
Modify a Policy 513
Delete a Policy 514
Apply Admission Control 515
Set Up UDP Latency Control 517
Set Up Rate Scaling 518
Set Up Web Redirection 520

Substitute Diffserv Values 522
Substituting a Diffserv Code Point Value 522
Substituting a Diffserv COS/TOS Value 523

Create VLAN Policies 524
Change the VLAN Priority (802.1p) 524
Change the VLAN Identification (802.1Q) 525

Enable/Disable Traffic Shaping 527

Graph and Report Tasks 528

Create a Traffic Assessment Report 529

View the Top Ten Pie Chart 535

List the Top Ten Classes 537

Create a Graph 538
Creating Graphs from the Top Ten Page 538
Creating Graphs from the Monitor Traffic Page 539
Creating Graphs from the Manage Tab 540
Analyzing Graphs 540

Create Top 10 Categories Graph 542

Graph a Specific Link, Partition, or Class 544
Additional Information *547*

Show Link Utilization 548

Graph Multiple Classes 550

Perform Host Analysis 552
Open the Host Analysis Tool 552
Host Analysis Table 552
Flow Analysis Table 555

Create Statistical Reports 557

Generate VoIP Reports 559

VoIP Latency Calculation 561

VoIP R-Factor and MOS 562
Measurement Statistics Tasks 563
View Measurement Statistics 564
Download Measurement Statistics to a File 567
Back Up Measurement Data 570
Restore Measurement Data 571
Reset Measurement Data 572
Host Accounting 573
Measure VoIP Statistics 574
VoIP Metrics Requirements *574*
Enabling VoIP Latency Calculations *574*
Viewing VoIP Metric Statistics *575*
Response Time Measurement Tasks 577
Monitor Response Times 578
Understanding the Monitor Response Time Window **578**
Show Transaction Delays 580
Set Thresholds 582
Total Delay Threshold *582*
Service Level Threshold *583*
Tracking Service Level Agreements *584*
View Response-Time Graphs 585
Enable Worst Client and Server Analysis 587
List the Worst Clients and Servers 589
Understanding "Worst" *590*
Example of Worst Server *590*
Adaptive Response Tasks 592
Enable Adaptive Response 593
Enable Adaptive Response Agents 594

Add a New Adaptive Response Agent 595
Edit an Adaptive Response Agent 596
Delete an Adaptive Response Agent 597
Temporarily Disable an Agent 597
Add or Edit Adaptive Response Action Files 598
PolicyCenter Action File Portfolios 598
Creating a New Action File 598
Action Output Files *599*
Editing an Action File 599
Reset Action File Parameters on an Agent 600
Enable or Disable Action Files 602
Distribute Action Files 603
View Drill-Down Incident Reports 604
Adaptive Response Agent Examples 606
Network Health 606
Unit Health 606
Application Health 606
Hosts 606
User Event Emulation 606
Create Partition Utilization Agents 607
Create a New Agent 607
Check the Agent Status 608
Create an Action File that Sends an Email Notification 609
Analyze Efficiency with Traffic Performance Agents 611
Create a Traffic Performance Agent 611
Check the Agent Status 612
Create an Action File that Sends an Email Notification 613
Use the Memory Allocation Agent 615
Create a Memory Allocation Agent 615
Check the Agent Status 616
Create an Action File 617
Use the System Load Agent 619

Create a System Load Agent 619
Check the Agent Status 620
Create an Action File 621

Use the Unit Limits Agent to Track Capacity 623
Create a Unit Limits Agent 623
Check the Agent Status 624
Create an Action File 625

Create Default Traffic Agents 627
Create a New Agent 627
Check the Agent Status 628
Create an Action File that Sends an Email Notification 629

Use the New Application Agent 632
Create an Agent Based on the New Application Template 632
Check the Agent Status 633
Create an Action File that Sends an Email Notification 635

Monitor Bandwidth of New Applications 637
Modify the Agent Settings 637
Check the Agent Status 638
Create an Action File 639

Identify and Control High Bandwidth Hosts 641
Create a High Bandwidth Host Agent 641
Create a Quota Bandwidth Host Agent 642
Check the Agent Status 643
Create an Action File 644

Use Agents to Detect Spoofing 646
Create an Exception Host List 646
Modify the Agent Settings 647
Check the Agent Status 648
Create an Action File that Sends an Email Notification 649
Restrict Bandwidth for Violating Hosts 650

Use Agents to Detect SYN Attacks 652
Create an Exception Host List 652
Modify the Syn Attack - Failed Flow Agent Settings 653
Check the Agent Status 654

Create an Action File that Sends an Email Notification 655
Restrict Bandwidth for Violating Hosts 656

Create an Agent that Monitors a Partition Variable 657
Create a New Agent 657
Check the Agent Status 658
Create an Action File that Sends an Email Notification 660

Create an Agent that Monitors a Class Variable 662
Create a New Agent 662
Check the Agent Status 663
Create an Action File that Sends an Email Notification 665

Create an Agent that Monitors a Link Variable 667
Create a New Agent 667
Check the Agent Status 668
Create an Action File that Sends an Email Notification 670

Event Notification Tasks 672

Set Up Event Notification 673
Defining the SNMP Trap Destination 673
Setting Up Email Notification 673
Setting Up Syslog on a PacketShaper 673

Set Up Email Notification 674
Define Email Recipients 674

Use Predefined Events 675
PacketShaper Predefined Event List 675

Define an Event 676

Register an Event 677

View Event Status 679

View Event Graphs 681

View Event Logs 683

Manage the Email Queue 685

Administration Tasks 686

Use the File Browser 687
Delete Files 688
Create a ZIP File 689
Open a ZIP File 691
Unzip a File 692
Download a File 693
Upload (Copy) a File to PacketShaper 694
Create a Command File 696
Review Log Files 699
Create a Schedule 700
List Schedules 704
Delete a Schedule 706
Edit a Schedule 707
Back Up a Configuration 708
Capture the Traffic Configuration 708
Save the Configuration 709
Back Up the Configuration Files 709
Restore a Configuration 710
Watch Mode Tasks 711
Configure Watch Mode 712
Select a Non-Inline Topology 713
Define Watch Mode Routers 714
Display Watch Mode Configuration 716
Delete Watch Mode Routers 718

Disable Watch Mode 719

Connect to a Network Tap 720

High Availability 721

Configure Standby 722

Configure PacketShapers for Standby 723

Verify Standby Partners 724

Standby Partner Failure 726

Returning to Active Status *728*

Disable Standby 730

Configure a High Availability Topology for Access-Link Monitoring 731

Define High Availability Access Routers 732

Define WAN Link Interfaces 734

Enable Access-Link Monitoring 736

View the High Availability Configuration 738

Modify High Availability Access Routers 740

Modify High Availability Interfaces 742

Remove High Availability Access Routers 743

Remove High Availability Interfaces 745

Customer Portal Tasks 747

Create Customer Portal Accounts 748

Edit Customer Information 750

Remove Customer Accounts 752

Set a System-Wide Message 753

Download Sample Portal Files 754

Customize the Login Screen 756
Design Customer Portal Pages 757
Assign a Customer Portal Address 758
Log Into the Customer Portal 759
Configure Captive Portal 760
Sky UI Tasks Step-by-Step 763
Sky UI Classification Tasks 764
Create a Traffic Class in the Sky UI 765
Exception Classes 768
Discovery within a Class 769
Create a Traffic Class with Multiple Criteria 770
Classify Users in the Sky UI 774
Create a User-Based Class in the Sky UI 775
Create a User Group Class in the Sky UI 777
Create a Class Based on a User List in the Sky UI 779
Open the User List Editor in the Sky UI 780
Create a User List in the Sky UI 782
Delete a User List in the Sky UI 784
Edit a User List in the Sky UI 785
Manually Create an IPv6 Class in the Sky UI 787
Classify URL Categories in the Sky UI 788
Create a Category-Based Class in the Sky UI 789
Auto-Discover Category Classes in the Sky UI 791

Discover Categories Under Web-Based Classes in the Sky UI 793
Create a Class for Multiple Categories in the Sky UI 794
Create a Web Application Class in the Sky UI 795
Create a VoIP Class in the Sky UI 797
Modify a Traffic Class in the Sky UI 799
Add Matching Rules in the Sky UI 800
Edit Matching Rules in the Sky UI 802
Use Host Lists in the Sky UI 804
Open the Host List Editor in the Sky UI 805
Create a Host List in the Sky UI 807
Edit a Host List in the Sky UI 809
Delete a Host List in the Sky UI 811
Sky UI Class Tree Management Tasks 812
Configure Class Tree Settings 813
Add a Folder in the Sky UI 816
Delete Classes in the Sky UI 818
Move Classes in the Sky UI 819
Copy Classes in the Sky UI 820
Rename a Class in the Sky UI 821
Create Class Lists in the Sky UI 822
Use Class Lists in the Sky UI 824
Modify a Class List in the Sky UI 825
Rename a Class List in the Sky UI 827

Copy the Class Tree to the Clipboard 828

Select Which Columns and Rows to Display in the Sky UI Class Tree 829
Hiding Statistics *829*
Sorting the Class Tree *830*
Filtering the Class Tree *830*

Update the Sky UI Class Tree Statistics 831

Reset the Sky UI Class Tree Statistics 832

Sky UI Traffic Control Tasks 833

Control Your Traffic 834

Create Partitions in the Sky UI 836

Set a Rate Policy in the Sky UI 838
Setting a Rate Policy *838*

Set a Priority Policy in the Sky UI 840
Setting a Priority Policy *840*
Priority Policy Examples *841*

Set a Never-Admit Policy in the Sky UI 842

Set a Discard Policy in the Sky UI 844

Set an Ignore Policy in the Sky UI 846

VoIP Optimization 848

DSCP Marking 850

Delete a Partition in the Sky UI 852

Remove a Policy in the Sky UI 853

Modify a Policy/Partition in the Sky UI 854

Sky UI Graphs & Reports Tasks 855

View the Sky UI Dashboard 856
Overview **856**
Adjusting the Graph Sizes **857**

Save the Sky UI Dashboard 858
Monitor Traffic in the Sky UI 859
Display Real-Time Graphs in the Sky UI 861
Save Graphs in the Sky UI 863
Display Historical Graphs in the Sky UI 864
Print Historical Graphs in the Sky UI 866
Save Historical Graph Data in the Sky UI 867
Reference - Look it up! 868
Policy/Partition Guidelines 869
Measurement Variables 873
Table of Common Variables 873
Table of Link Variables 875
Table of Partition Variables 877
Table of Class Variables 878
Applications, Protocols, and Services Classified by PacketShaper 11.9 884
Service Groups 884
Services Table 885
URL Categories 907
Preconfigured Graphs (Advanced UI) 919
Graphs for Analyzing Bandwidth Usage 919
Graphs to Analyze Efficiency 920
Graphs for Analyzing Top Ten 920
Graphs for Analyzing Response Time 920
Graphs for Analyzing TCP Connections 921
Sky UI Graphs 922
PacketShaper Port and URL Usage 924
Glossary of Terms 926
Tutorials 933

PacketShaper Feature Matrix 934

Product Information 936
Log in to PacketShaper Browser Interface 936
Log Out of PacketShaper 937

First Steps to Using... 938

PacketShaper S-Series Product Specifications 939

PS-S400 Product Specifications 940

PS-S400 Device and Port Names 942

PS-S400 NIC Support 943

PS-S500 Product Specifications 944

PS-S500 Device and Port Names 946

PS-S500 NIC Support 947

PacketShaper S-Series Network Interface Capabilities 949

PacketShaper S-Series Configuration Limits 950

PacketShaper S200 Configuration Limits 951

PacketShaper S400 Configuration Limits 953

PacketShaper S500 Configuration Limits 954

Power 955

PS-S200 Power 956
Powering On 956
Power Cycling 956
Powering Off 956

PS-S400 Power 957
Powering On 957
Power Cycling 957
Powering Off 957

PS-S500 Power 959

Powering On 959
Power Cycling 959
Powering Off 960

Deployment Topologies 961

Distributed Branch Topology 963
Environmental Requirements 963
Reasons to Choose 963
Installation Considerations 964
Capabilities 964
Limitations 964
Configuration Suggestions 964
Topology Example 965
Scaling Considerations 965
Configuration Suggestions 966

Topologies with Firewalls 967
Installation Considerations 967
Capabilities 968
Limitations 968
Configuration Suggestions 968
Scaling Considerations 968

Core Site Internet Link Topology 969
Environmental Requirements 969
Reasons to Choose 969
Installation Considerations 969
Capabilities 970
Limitations 970
Configuration Suggestions 970
Scaling Considerations 971

Core Site WAN Link Topology 972
Environmental Requirements 972
Reasons to Choose 972
Installation Considerations 973
Traffic Tree Options 973

Topologies with Multiple LANs 974
Installation Considerations 974

Capabilities 974
Topology Examples with Multiple NICs 975
Limitations 975
Configuration Suggestions 976
Scaling Considerations 976

High Availability Topologies 977
Installation Considerations 977
Topology Examples 977
Basic Failover Topology *978*
Load Balancing Topology *979*
Load Balancing with Multiple VLANs *980*
Topologies with Firewalls *982*
Capabilities 984
Limitations 984
Configuration Suggestions 984
Scaling Considerations 984

Traffic Tree Options 985
1. Application-Based Traffic Tree 986
2. Simple Location-Based Tree with Locations Only 986
3. Location-Based Tree with Per-Location Applications 988
4. Location-Based Tree with Application Groups 989
5. Location-Based Tree with Global Applications 990

Virtual Private Network Topology 993
Environmental Requirements 993
Reasons to Choose 993
Installation Considerations 993
Capabilities 994
Configuration Considerations 994
Scaling Considerations 994

Non-Inline Deployments (Watch Mode) 996
Reasons to Choose 996
Installation Considerations 996
Topology Examples 996
Example 1: Monitor One LAN *996*
Example 2: Monitor Multiple LANs *997*
Additional Non-Inline Topology Examples *998*
Capabilities 999
Limitations 999

Configuration Suggestions 999
Scaling Considerations 1000

Access the CLI with a Browser 1001

Access the CLI Using a Console Connection 1002

Access the CLI Using Remote Login 1003

Access the CLI with the Quick CLI Commands Utility 1004
Types of Commands 1004
Selecting Multiple Classes 1004
Running a Quick Command on a Group of Classes 1005
Running other CLI Commands 1005

Access the Command-Line Interface 1006

Agent Host Lists 1007
Host Exception Lists 1007
Violating Host Lists 1007
Restricting Bandwidth for Violating Hosts 1008

Application-Specific Matching Rule Criteria 1009
Overview 1009
Citrix-ICA Criteria 1009
Classifying by Published Application 1010
Classifying by Client Name 1011
Classifying by Priority 1011
DCOM Criteria 1011
DICOM Criteria 1012
FTP-Data-Clear and FTP-Data-Clear-IPv6 Criteria 1012
HTTP, HTTP-IPv6, and SOAP-HTTP Criteria 1013
Host DNS Name or IP Address Classification 1013
URL Classification 1013
Content Type Classification 1014
Web Browser (User Agent) Classification 1014
Understanding Web Criterion Classification 1015
HTTP-Tunnel and HTTP-Tunnel-IPv6 Criteria 1016
ICMP Criteria 1016
NNTP-Clear Criteria 1017
Oracle-netv2 Criteria 1017
Automatic Creation of Database Classes 1017
Manual Creation of Database Classes 1017
PostgreSQL Criteria 1018

Automatic Creation of Database Classes *1018*
Manual Creation of Database Classes *1018*
ProxySG-ADN and ProxySG-ADN-IPv6 Criteria 1018
Automatic Creation of ProxySG-ADN Application/Port Classes *1018*
Manual Creation of ProxySG-ADN Application/Port Classes *1019*
RTCP-I Criteria 1019
RTP-I Criteria 1020
SMTP-Clear and SMTP-Clear-IPv6 Criteria 1021
SSL and SSL-IPv6 Criteria 1022
WAP 1022
Specifying Multiple Criteria 1023

Bandwidth Allocation 1024

BCAAA Troubleshooting 1025
Troubleshooting Tools 1025
Problem: Can't connect to the BCAAA server 1026
Problem: Can connect to BCAAA server, but user awareness isn't functioning on PacketShaper 1027

Browser Settings 1028

Examples of Bandwidth-Utilization Analysis 1029

Classification Hints and Examples 1031

Command Line Overview 1033
Command Usage Conventions 1033
Typographical Conventions 1033
Editing Previously Entered Commands 1034

Command Line Overview 1035
Look vs. Touch Access *1035*
Command Usage Conventions *1035*
Typographical Conventions *1035*
Editing Previously Entered Commands *1036*

Create a Host List 1037

Device Matching Rules in PolicyCenter 1039

Device 1040

Examples of Efficiency Analysis 1041

FDR Notes 1042

FDR Troubleshooting 1043

Verify Settings 1043

Test Network Connectivity 1043

Verify that FDRs are Being Sent 1043

Host Agent Templates 1045

Network Health Agent Templates 1057

Unit Health Agent Templates 1059

Application Health Agent Templates 1062

User Event Emulation Agent Templates 1065

Access Guided Setup With a Console Connection 1067

Info Tab in the Sky UI 1068

Policy Inheritance Rules 1069

Location-Based Traffic Tree Example 1071

View the Operations Log in the Sky UI 1072

Look Up ifName and ifIndex Values 1073

Matching Rule Details 1074

Overview *1074*

Protocol *1074*

Server Location *1074*

Hosts *1074*

Service Groups *1076*

Services *1077*

URL Categories *1077*

Users *1078*

1078

Specific Sub-Protocol *1078*

Web-Specific Parameters *1079*

Application-Specific Criteria *1079*

Device *1080*

IP Precedence (Diffserv) *1080*

Virtual LAN Identification (802.1Q and ISL) *1081*

VLAN Priority (802.1p) *1082*

Modify a Host List ... 1083
NetFlow v5 Header ... 1084
NetFlow v5 Record Format ... 1085
PacketShaper Folders ... 1087
PacketShaper SNMP Traps ... 1088
Arithmetic Comparisons ... 1090
Arithmetic Example ... 1090
Select Traffic Classes for the Portal Page ... 1091
Search Facility ... 1091
Using the Same Portal Page on Multiple Units ... *1092*
Using Different Portal Pages on a Unit ... *1092*
Display Custom Information on the Portal Page ... 1093
Include Graphs on a Portal Page ... 1094
Saving the Web Page ... 1094
Link to Other HTML Pages ... 1096
Display Statistics on the Portal Page ... 1097
Error Detection ... *1097*
Measurement Engine Data ... *1097*
Displaying Current Class Name ... 1098
Meta-data ... *1099*
Customer Portal API ... 1100
Measurement Variables Available to Portal Pages ... 1103
Pre-Installation Check List ... 1104
Priority Guidelines ... 1106
Protocol Family ... 1107
ProxySG Optimization ... 1108
RADIUS Accounting Troubleshooting ... 1109
PING the Accounting Server ... 1109
Send Accounting Test Messages ... 1109
Checking the Accounting Log ... 1110

Checking the Accounting Server Status 1111

RADIUS Authentication Troubleshooting 1112
PING the Authentication Server 1112
Send a Test Login Request 1112
Checking the Authentication Server Status 1113

Rate Policy Examples 1114

Response Time Measurement's Technical Details 1115
Determining the Round Trip Time (RTT) 1115
Determining the Server Delay 1116
Determining the Total Delay 1116
Determining the Network Delay 1117

Troubleshooting Response Time Measurement Issues 1118
Accessing Variables After Upgrades 1118
Data for Manually Created Classes 1118
Browser Data Differs From Command-Line Data 1118
Graph Issues 1119

Sample Customer Portal Page with 60 Minute Graphs 1120

Sample Customer Portal Page 1121

Secure Logins 1127
Browser Interface 1127
Logging in Securely the First Time 1127
Command-Line Interface 1127

Windows SecureCRT Settings 1128

Security Alert 1129

Selecting a Time Period 1130

Service Groups Best Practices 1131
Initial Setup of Service Groups 1131
Deleting Custom Groups 1131
Service Group Maximum 1132
Browser Refresh may be Required 1132

SNMP Trap Filtering 1133

Supported Ciphers and Algorithms for SSH and SSL 1134
SSH *1134*
SSL *1135*

PacketShaper Syslog Warn Messages 1136

Characterizing Traffic 1137

Traffic Tree Designs 1139
Application-Based Tree 1140
Capabilities *1141*
Limitations *1141*
Scaling Considerations *1141*
Simple Location-Based Tree with Locations Only 1141
Capabilities *1141*
Limitations *1142*
Complementary Products *1142*
Scaling Considerations *1142*
Configuration Suggestions *1142*
Location-Based Tree with Per-Location Applications 1143
Capabilities *1143*
Limitations *1144*
Scaling Considerations *1144*
Configuration Suggestions *1144*
Location-Based Tree with Global Applications 1145
Capabilities *1146*
Limitations *1146*
Scaling Considerations *1146*
Inheritance *1146*
Configuration Suggestions *1147*

URL Syntax 1148

Use PacketShaper MIBs 1150
actionfile sync 1152
agent action 1153
agent createdefaults 1154
agent delete 1154
agent interval 1154
agent new 1155
agent off 1156
agent on 1157
agent override 1157
agent parm 1157
agent show 1158

arp 1162
banner show 1163
cat 1163
cd 1164
class capture-ids 1164
class category 1164
class copy 1165
class criteria 1165
class criteria attributes 1168
class criteria recent 1170
class criteria track 1170
class delete 1171
class discover 1171
class group delete 1171
class group move 1172
class group new 1173
class group override 1173
class group reinherit 1173
class group reset 1174
class group show 1175
class hosts 1175
class id 1176
class licenses 1176
class load 1177
class move 1177
class new 1177
class note 1179
class owner 1179
class rename 1180
class reset 1180
class rule 1181
class services 1182
class set 1182
class show 1183
class test 1185
class user-group 1186
class user-services delete 1187
class user-services new 1187
class user-services show 1189

class users 1189
class web-app disable 1190
PolicyCenter Support *1190*
class web-app enable 1191
class web-app show 1191
config clear 1192
config dump 1192
config errors 1192
config load 1192
config reset 1193
config save 1194
config secure 1195
config setup 1195
config unset 1195
date 1196
dns lookup 1196
dns names 1196
dns refresh 1197
dns rlookup 1198
dns servers 1198
dns trace 1198
email queue 1198
email retry 1199
email test 1199
event delete 1199
event email 1199
event log reset 1200
event log status 1200
event new 1200
event register 1202
event reset 1203
event show 1203
event test-email 1203
event unregister 1204
exit 1204
help 1204
highav add 1204
highav community 1205
highav delete 1205

highav disable 1205
highav enable advanced 1206
highav enable basic 1206
highav interface add 1206
highav interface delete 1207
highav interface modify 1207
highav override 1208
highav show 1209
history 1210
hl add 1211
hl delete 1211
hl new 1212
hl override 1213
hl refresh 1213
hl resolve 1213
hl rm 1214
hl show 1214
host accounting categories 1214
host accounting enable 1215
host accounting retrieve 1216
host accounting show 1217
hostdb cache 1218
hostdb info 1219
hostdb rtostats 1221
hostdb show 1222
hostdb side auto 1224
hostdb side default 1224
hostdb side manual 1224
hostdb side reset 1225
hostdb side rm 1225
hostdb side set 1226
hostdb side show 1227
hostdb topusers 1228
ipfilter 1230
ipfilter clear 1232
ipfilter iponly 1233
ipfilter show 1233
license birthcert 1234
license load 1234

license update 1235
links show 1235
ls 1236
measure cumvar 1236
measure dump 1237
measure dump by element 1237
measure dump by event 1239
measure dump by time 1241
measure dump by var 1244
Measure Dump Examples 1246
Measure Dump "by var" and "by element" Comparison 1247
measure reset 1248
measure show 1249
measure start 1251
measure stop 1251
mib 1252
more 1252
net nic 1252
packetcapture add 1253
packetcapture limit packets 1254
packetcapture off 1255
packetcapture on 1255
packetcapture remove 1256
packetcapture status 1258
partition apply 1258
partition dynamic apply 1260
partition dynamic cap 1261
partition dynamic cleanup 1262
partition dynamic remove 1262
partition dynamic summary 1263
partition remove 1263
partition show 1263
ping 1266
policy admit 1267
policy apply discard 1267
policy apply ignore 1268
policy apply never-admit 1268
policy apply priority 1268
policy apply rate 1269

policy default 1270
policy delaybound 1270
policy dscp 1271
policy excess 1271
policy flowlimit 1271
policy guaranteed 1272
policy precedence 1273
policy remove 1273
policy route 1273
policy show 1273
policy substitute 1274
policy test 1274
policy tos 1275
policy vlan 1275
portal delete 1276
portal modify 1277
portal new 1278
portal show 1279
pwd 1279
radius acct 1279
radius chaplogin 1280
radius clear 1281
radius login 1281
radius session 1283
radius show 1284
reset 1284
rtm accept 1285
rtm clear 1285
rtm drilldown 1285
rtm hosts 1286
rtm show 1286
rtm threshold 1286
rtm worst 1287
run 1287
schedule delete 1288
schedule disable 1288
schedule enable 1288
schedule new 1289
schedule show 1292

send email 1292
send syslog 1293
send trap 1293
setup access default 1294
setup access disable 1294
setup access enable 1295
setup access show 1295
setup adaptiveresponse 1296
setup bcaaa cache 1296
setup bcaaa force-register 1297
setup bcaaa server-test 1298
setup bcaaa service 1299
setup bcaaa show 1300
setup bcaaa timeout 1301
setup bypass 1301
setup capture 1302
setup date 1303
setup discover 1303
setup dns 1303
setup domain 1304
setup email 1304
setup flowrecords engineID/engineType 1305
setup flowrecords filters 1305
setup flowrecords id 1307
setup flowrecords show 1308
setup gateway 1309
setup guide 1309
setup heartbeat 1309
setup images active 1310
setup images add 1311
setup images detail 1311
setup images lock 1312
setup images remove 1312
setup images show 1313
setup images unlock 1314
setup ipaddress 1314
setup keys show 1314
setup link 1315
Considerations *1315*

setup load 1316
setup loadshedding 1316
setup message 1319
setup nic 1320
setup password 1321
setup portal ip 1322
setup portal message 1322
setup portal show 1322
setup radius acct 1322
setup radius auth 1323
setup radius interval 1324
setup radius limit 1325
setup radius method 1325
setup radius show 1326
setup reset 1326
setup secure 1327
setup shaping 1328
setup show 1328
setup snmp accessgroup 1329
setup snmp destinations 1331
setup snmp disable 1332
setup snmp enable 1332
setup snmp look|touch 1333
setup snmp remoteuser 1333
setup snmp show 1335
setup snmp syscontact 1336
setup snmp syslocation 1337
setup snmp sysname 1337
setup snmp target 1338
setup snmp user 1340
setup snmp view 1342
setup sntp 1344

setup ssl install 1345
setup standby 1345
setup standby partner 1346
setup syslog add 1347
setup syslog rate 1348
setup syslog remove 1349

setup syslog show 1349
setup syslog state 1350
setup tacacs acct 1350
setup tacacs auth 1351
setup tacacs clear 1352
setup tacacs method 1352
setup tacacs show 1353
setup tacacs timeout 1353
setup timezone 1354
setup variable 1355
setup web-proxy 1362
setup web-proxy default 1362
setup web-proxy server 1363
setup web-proxy show 1363
setup webpulse 1364
setup webpulse cache-test 1364
setup webpulse cache-update 1365
setup webpulse discovery 1366
setup webpulse express 1367
setup webpulse map-download 1368
setup webpulse reset 1369
setup webpulse show 1369
shutdown 1371
synthetic add 1372
synthetic delete 1373
synthetic options 1373
synthetic show 1374
sys 1375
sys limits 1375

tenant add customer 1377

tenant add customer aup 1378

tenant add site 1379

tenant add site <attribute> 1380

tenant enable 1382

tenant remove 1383

tenant set 1384

tenant show 1385

traffic active 1385
traffic bandwidth 1386
traffic flow 1388
traffic guaranteed 1392
traffic history 1392
traffic licenses 1394
traffic statistics 1394
traffic tree 1395
ul add 1396
ul delete 1397
ul new 1397
ul override 1398
ul reformat 1398
ul rm 1399
ul show 1399
uptime 1399
version 1400
watch add 1400
watch delete 1400
watch show 1401

Bytes Transmitted Graph 1402

Connection Retransmissions Graph 1403

Guaranteed Rate Failures Graph 1404

Network Efficiency Graph 1405

Normalized Network Delay Graph 1406

Packet Round Trip Time Graph 1407

Packets Transmitted Graph 1408

TCP Connections Initiated Graph 1409

TCP Health Graph 1410

Transaction Delay Graph 1411
Average Utilization Graph 1412
Peak Utilization Graph 1413
Average Transaction Size Graph 1414
Bytes Transmitted Graph 1415
Class Utilization with Peaks Graph 1416
Class Utilization Graph 1417
Connection Retransmissions Graph 1418
Dynamic Partition Usage 1419
Guaranteed Rate Failures Graph 1420
Link Utilization with Peaks Graph 1421
Link Utilization Graph 1422
Network Delay Distribution 1423
Network Delay Graph 1424
Network Efficiency Graph 1425
Normalized Network Delay Graph 1426
Packet Round-Trip Time Graph 1427
Packet Size Distribution Graph 1428
Packets Transmitted Graph 1429
Partition Utilization with Peaks Graph 1430
Partition Utilization Graph 1432
Server Delay Distribution Graph 1433
Server Delay Graph 1434

Service Level Compliance Graph 1435
TCP Connections Initiated Graph 1436
TCP Health Graph 1437
Top 10 Children Classes Graph 1438
Top 10 Classes Graph 1439
Top 10 Children Classes vs. Top 10 Classes Graphs 1440
Top-10 Partitions Graph 1441
Transaction Delay Distribution Graph 1442
Transaction Delay Graph 1443
Allow URLs within a Blocked Category 1444
Allow Users Access to a Blocked Category 1445
Block URLs within an Allowed Category 1446
Block Users Access to an Allowed Category 1447
Remove Units from PolicyCenter 1448
Switch the Unit to Local Mode 1448
Delete the Unit Configuration from PolicyCenter 1448
Print and Update an Event Graph 1449
Set Up Syslog 1451
Define the Syslog Server 1451
Facility Types 1451
Severity Levels 1452
Enable the Logging Feature 1453
Update the Event Statistics 1454
Disable Access-Link Monitoring 1455
Define Host Accounting Categories 1456
Enable Host Accounting 1457

Retrieve Host Accounting Data 1458

Create Custom Reports 1459

Additional Resources 1461

Create Standard Reports 1462

Additional Resources 1464

Delete Reports 1465

IPv6 Traffic Summary Report 1466

Select the Time Period 1468

Select a Different View 1470

Getting More Details About Top Classes 1471

Print and Update a Custom Report 1472

Retrieve Saved Reports 1473

Save Graphs as GIF Files 1474

Saving in Microsoft Internet Explorer 1474

Saving in Firefox 1474

Save Reports as PDF Files 1475

Display Real-Time Graphs 1476

Additional Resources 1476

Change the Time Period 1477

Reset RTM Statistics 1478

Update the Response Time Statistics 1479

Using this Help System 1480

Find the Information you Need 1480

Printing a Topic 1480

Glossary 1481

Overviews - The Big Picture

High-level, narrative descriptions of PacketShaper products, their features, and other product topics that merit extra exploration

Access-Link Monitoring Overview

Redundant network configurations typically involve some type of load-balancing or load-sharing scheme that determines how traffic is distributed across the available WAN links. In some configurations, the load-balancing scheme may be unable to enforce distribution of traffic so that each available WAN link is utilized 100 percent, nor can it ensure that no links will ever be overloaded. In addition, there is always the potential that any given link or router could go down, reducing the total available capacity to the WAN links remaining.

PacketShaper's access-link monitoring feature allows PacketShaper to deal with this "imperfect" load-balancing issue and has the ability to respond to the occurrence of WAN link failure. When access-link monitoring is enabled, PacketShaper can adjust partitions appropriately to prevent overloading any given WAN link and to account for lost available capacity due to router or link failure. Access-link monitoring has two modes: basic and advanced.

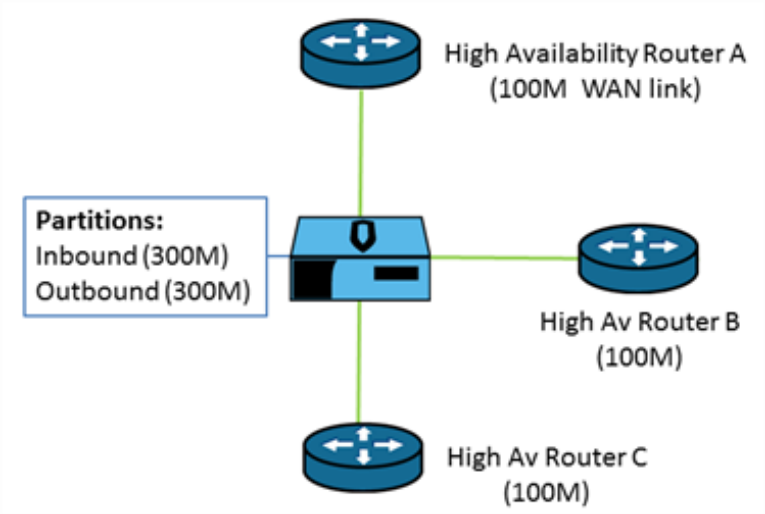

When access-link monitoring is configured in the **basic mode** (shown above), the PacketShaper polls the configured router(s) every 30 seconds to assess the status (link up or link down) of the WAN link interfaces. If a link goes down, PacketShaper will automatically adjust the total available capacity by subtracting out the capacity of the down link. As part of this process, it will adjust the access link size and resize Inbound and/or Outbound partitions to reflect the available bandwidth. This situation is illustrated in the diagram below.

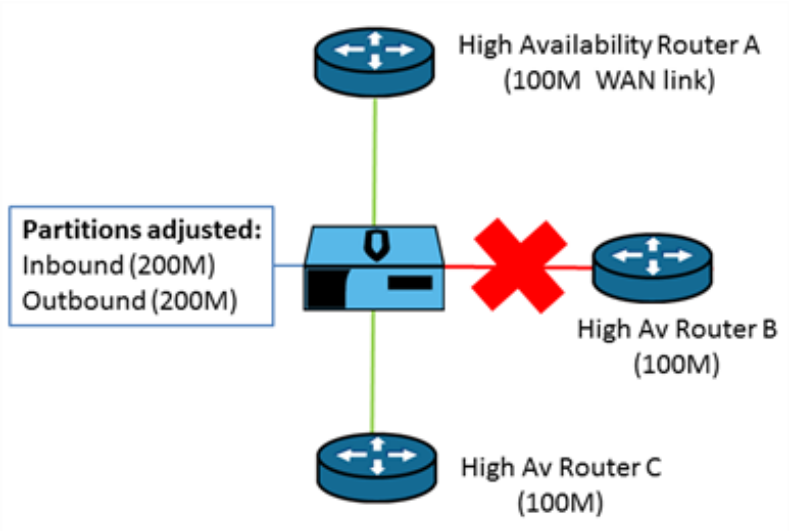

When **advanced mode** is enabled, PacketShaper can help prevent the overloading of an interface. The PacketShaper will use SNMP polling to assess the actual throughput of each configured WAN link interface; the configured routers are polled every 30 seconds. When an interface approaches 25% of its configured capacity, PacketShaper will begin pacing the traffic sent to the router to prevent overloading any interface. This pacing will also greatly reduce the number of retransmissions. PacketShaper begins adjusting the partition sizes early in order to ensure gradual, smooth adjustments, as well as to give you time to modify policies if desired. PacketShaper will poll the router frequently, and once there is evidence that the links are out of danger of being overloaded, it will gradually increase the size of the partition(s).

Access-Link Monitoring Requirements

The access-link monitoring feature has the following requirements and limitations:

- In order to have the ability to adjust partition sizes (a critical part of the access-link monitoring feature), you must enable traffic shaping on your PacketShaper.
- SNMP must be enabled on each high availability router.
- Although there is no pre-set limit for the number of routers each PacketShaper can monitor, the polling process consumes both CPU and memory resources. For best access-link monitoring performance, make sure that your unit is not already operating at its maximum capacity.

MIB Variables Polled by the Access-Link Monitoring Feature

As mentioned previously, the access-link monitoring feature uses SNMP polling to assess the interface status and throughput. The following MIB variables are requested in both basic and advanced mode:

- sysName– name that identifies the router
- ifName–name that identifies the interface
- ifOperStatus–indicates if the given interface is up or down
- ifSpeed–link speed

In advanced mode, the following additional MIB variables are polled:

- ifInOctets–number of Inbound bytes of traffic seen on the interface
- ifOutOctets–number of Outbound bytes of traffic seen on the interface

See Configure a High Availability Topology for details on configuring access-link monitoring.

Adaptive Response Overview

This overview gives a general description of the adaptive response feature. For more detailed overviews on specific components of this feature, refer to the Adaptive Response Agents Overview and Adaptive Response Action Files Overview.

*For **examples** and **workflows** showing how to use this feature to Monitor and Respond to Conditions of Interest on your network, see* Monitor and Respond to Conditions of Interest using Agent Templates and Monitor and Respond to My Own Custom Condition.

Introduction

Network problems. Hardware problems. You never know when they are going to strike, but they almost always seem to come at an inconvenient time. Even if you have someone monitoring the health of your entire network 24/7, very often by the time a problem has become critical, users have already been impacted. PacketShaper's adaptive response feature makes it easier to monitor your network by enabling an individual PacketShaper to monitor unit, application, and network health, hosts, traffic classes, links, and partitions, and by providing a color-coded summary that lets you identify potential problems at a glance.

The adaptive response feature has three major components: *agents*, *incident reports*, and *action files*.

Each performance standard is measured by an *agent*. You can create an agent that monitors the efficiency of just one mission-critical traffic class, for example, or measure the traffic performance of your entire network. Because acceptable behavior on one network or unit might be considered very unacceptable on another, you can set the thresholds of each agent to reflect your own definitions of good and bad performance. Each PacketShaper unit or PolicyCenter configuration can have a maximum of 32 agents.

PacketShaper comes with several different agents already configured. As soon as you enable an agent, adaptive response immediately starts monitoring your unit or network traffic, returning values and measuring them against their predefined good and bad performance thresholds. If there is a problem, the agent will send information about the problem to two different places: an incident report and an action file.

Incident reports show supporting data for the problem at the time it occurred, letting you see exactly what just happened. Incident reports are automatically created by the agent, and do not need to be created or configured by the user.

Action files are user-created scripts of CLI commands that are triggered to run when unit or network performance starts to degrade, or to rearm when the situation starts to return to normal. Very often, it is not enough to merely identify a problem area in an incident report. Once a problem has been pinpointed, corrections and adjustments still need to be made to solve the immediate problem. These customized action files can make automatic changes based on the problem at hand, often correcting the problem before users are ever impacted. Action file scripts can be configured to send notifications via email, syslog, and SNMP traps. Once the unit or network returns to normal, a second action file script can undo the previous changes or run additional commands.

Adaptive Response Agents

An agent is an entity in adaptive response that measures and monitors the performance of the PacketShaper, the network, or a specific class. A number of agents are predefined for you and will begin measuring and monitoring as soon as you have enabled

the agents you want to use. You can create additional agents using the included templates. (A template is a form that is used as a guide for creating agents.)

Adaptive response has five categories of agents:

- **Unit Health agents** monitor capacities and loads on the PacketShaper.
- **Network Health agents** monitor the efficiency of a specified traffic class, and how much of a partition is currently utilized.
- **Application Health agents** monitor the new traffic on the network, as well as traffic that is classified into */default* categories
- **Host agents** detect hosts that are using up too much of the link or identify hosts that may be attacking your network.
- **User Event Emulation agents** monitor any of the 130+ PacketShaper measurement variables.

The adaptive response feature includes agent templates for creating and defining agents, as listed in the following table. Click any of the agent template names below for detailed information.

Unit Health Agent Templates	Network Health Agent Templates	Application Health Agent Templates	Host Agent Templates	User Event Emulation Agent Templates
Unit Limits	Traffic Performance	Default Traffic	High Bandwidth Host	Class ME Variables
System Load	Partition Utilization	New Application	Host Info Variables	Partition ME Variables
Memory Allocation		High Bandwidth New App	Failed Flow Ratio	Link ME Variables
			New Flows Per Minute (NFPM) Failed Flow	
			New Flows Per Minute (NFPM) Side Unknown	

The following agents are predefined:

Agent Name	Based on this Template
High Bandwidth New App	High Bandwidth New App
Inbound Default Traffic	Default Traffic
Outbound Default Traffic	Default Traffic
Inbound Packet Drops	Link ME Variables
Outbound Packet Drops	Link ME Variables
Spoofing-Client	NFPM Side Unknown
Spoofing-Server	NFPM Side Unknown
Syn Attack-Failed Flows	Host Info Variables

> **Note:** Some agents are not predefined and must be individually configured if you want to use them.

You can enable or modify the settings of the default agents, create agents based on other templates, or create multiple instances of an agent based on different parameters or variables (not applicable to all agents). The **adaptive response** setup page allows you to create new agents; enable, edit, or delete existing agents; and enable or disable action files.

Adaptive Response Dashboard

The Info tab has five colored status indicators that show the status of each of the five categories of agents.

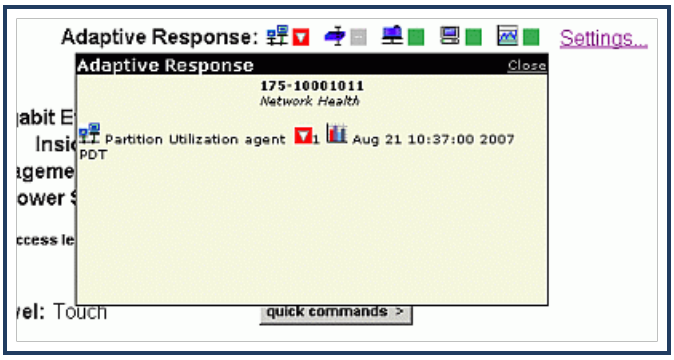

The category status color is based on the status color of the agents in that category.

Category Status Color	What it Means
Green	All agents in the category are reporting desired performance
Yellow	At least one agent in the category is reporting marginal performance
Red	At least one agent in the category is reporting unacceptable performance
Gray	All agents in the category are either disabled, absent, or have not yet reported a value. A status indicator will be gray until the agents in the category have had time to gather and analyze data. Another reason for a gray status is if there are no active agents in that category– either none are configured or all of the configured agents have been temporarily turned off.
Blue	At least one agent in the category has not been able to measure its target or return a value; when an agent has a blue status it indicates that the agent had a problem retrieving data. For example, an agent was set to measure a traffic class, but the class was subsequently deleted from the traffic tree.

Are all five status indicators showing green? You don t need to look any further to know that every agent in each category is reporting desired performance for the part of the network it is monitoring. Did a green status indicator just turn yellow? You can quickly find out exactly which particular agent in that category is reporting a problem—simply hover your mouse over the agent category to display a pop-up window with a separate status indicator for each individual agent in that category.

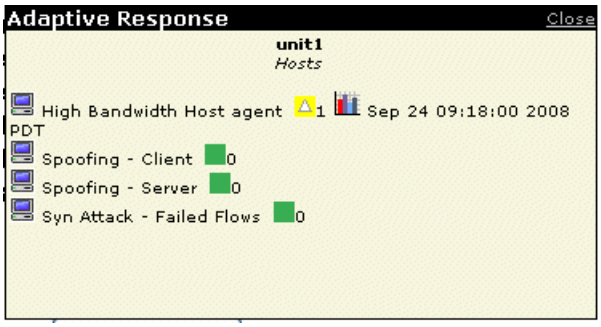

These pop-up windows can also contain links to additional valuable information. When an agent changes status from green/yellow to red, it will create a drill-down report (an incident report) that shows related data or a graph at the time the problem occurred. If a report is generated, a link to the report appears beside the status indicator for the agent. Some agents also create incident reports when status goes from red/yellow to green.

The colored status indicator for the category as a whole always reflects the agent in that category with the highest severity level, as shown below.

Highest Severity <– red – yellow –blue– green– >Lowest Severity

As an example, suppose you had two Unit Health agents with a green status, and one with a yellow status. The highest severity level for those agents would be yellow, and the Unit Health category would display a yellow status indicator on the Info tab. If one of the agents turned from green to red, the highest severity level would be red, and Unit Health would display a red status indicator.

Drill-Down Incident Reports

An incident report is a table or graph that adaptive response automatically generates when an agent changes status. A red incident report shows supporting data for the interval in which the agent crossed the red threshold (the time that the problem occurred). A green incident report shows the data for the interval in which the agent returned to a green status.

Adaptive response generates a red incident report if the agent's very first evaluation interval records a red status or when the status changes from:

- green to red
- yellow to red
- blue to red

In addition, some agents generate a green incident report when the status changes from:

- red to green
- yellow to green

The figure below shows an example incident report for a High Bandwidth Host agent. The incident report lists all the hosts for the interval, with the host using the most bandwidth at the top of the list.

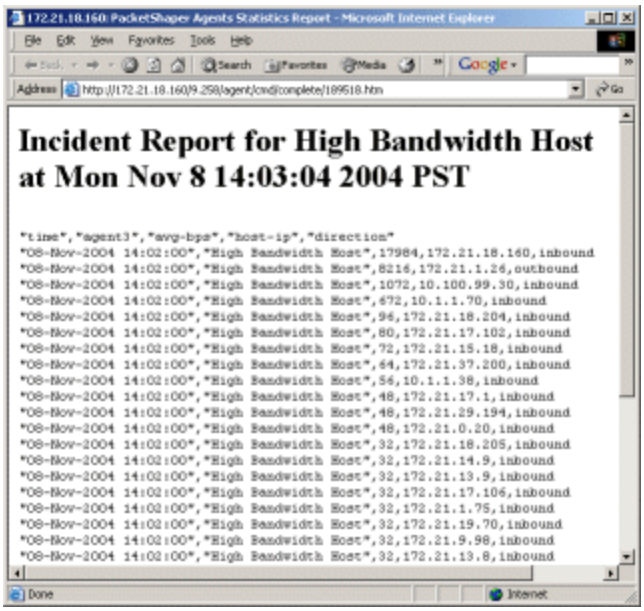

You may view incident reports from the browser interface of the PacketShaper or from PolicyCenter.

Adaptive Response Action Files Overview

Adaptive response action files are command files that are associated with a specific agent. Action files can be configured to take corrective action, notify network administrators of changes in their network, or to capture relevant history when an agent crosses a threshold to get worse or better. An action file can contain any CLI command that can run without additional user input. For example, the CLI command **setup discover on** is a useful command in an action file for a Default Traffic agent. When too high a percentage of traffic falls into the target /default class, the Default Traffic agent triggers its red action file. The action file issues the **setup discover on** command which turns on traffic discovery, improving the classification of traffic that was otherwise falling into the /default class. Action files can also include commands for sending notification that an agent has crossed a threshold and variables, which can be user-defined or automatically defined by the agent itself.

Each agent can have two assigned action files, one for the red threshold, and one for its green threshold. These files are triggered when any of the following thresholds are crossed:

- Red to Green–triggers green action file
- Yellow to Green–triggers green action file
- Green to Red–triggers red action file
- Yellow to Red–triggers red action file
- Blue (error state) to Red–triggers red action file

A red action file will also trigger if the first value returned by a new agent is past the red threshold. (This avoids the problem of the red action file never triggering if the agent always stays within the red threshold.) An action file gets triggered only when an agent changes status (from green/yellow to red, from red/yellow to green). For example, suppose an agent goes from green to red in a particular evaluation interval. The red action file will get executed. If the agent is still red in the next evaluation interval, the red action file will NOT get triggered. It's not until the agent goes to green or yellow and then back to red that the red action file gets triggered again.

Here is an example of a red action file for the High Bandwidth Host agent. This action file gets triggered when a single host has exceeded a certain percentage of the link size (the red threshold).

```
class new $direction $host-ip outside host:$host-ip
 class rule add $direction/$host-ip inside host:$host-ip
 send email raltman@test.com "High Bandwidth Host found" "Host $host-ip recently used $avg-
bps bps, which is more than the Red Threshold ($RedThreshold percent). Class
$direction/$host-ip was created."
```

This action file has three actions: creates a class based on the IP address of the bandwidth abuser, adds a matching rule to this class, and sends an email message that notifies the administrator of this situation. Note that this command file makes use of several action file variables ($host-ip, $direction, $avg-bps, and $RedThreshold). See Action File Variables below for more information.

Including Automatic Notification in an Action File

If you want to be notified automatically when an agent has crossed a red or green threshold, you can add one of the **send** CLI commands to the red or green action file. The notification can take one of the following forms: syslog entries, SNMP traps, or

email messages. For additional information on setting up and using these commands, click any of the links below.

- send email
- send syslog
- send trap

Creating an Action File

Action files must be stored in your PacketShaper's *9.258/agent/cmd* directory. Although you can use any text editor (such as Notepad) to create the action file, you may find it more convenient to use PacketShaper's File Browser utility which has a facility for creating command files and storing them directly on the unit. To access the File Browser utility, use one of the following ways:

- The **file browser** button on the Info tab
- The **browse files** button on the *Edit Agent* page

The second method offers two advantages. First, it automatically stores the action file in the proper place. (You don't have to navigate to the 9.258/agent/cmd directory). Second, after creating the action file, you are on the right screen for supplying the red or green action filename.

For step-by-step instructions on creating an action file, see Add or Edit Adaptive Response Action Files.

Action File Variables

By including variables in an action file, you can create action files that work in a variety of situations, that are applicable to multiple agents. Three types of variables can be included in action files:

- Variables available to all agents
- Agent-specific variables
- User-defined variables

Variables Available to All Agents

When the following action file variable names are included in an action file, the agent automatically enters the values for the variables as the file is run. You cannot change the values of these variables; their values come from the agent only.

Variable	Description
$agentname	Name of the agent that crossed a threshold
$endtime	The time at the end of the evaluation interval. For example, if an agent has 1-hour evaluation interval and starts measuring data at 8:00 AM, even if a threshold was crossed at 8:25, the agent s reported endtime will be 9:00 AM, the end of that interval.

Variable	Description
$scorevalue	The value of the score at the end of the interval, as measured by the agent
$scorecolor	Integer representing the status color of the agent, where 0=green, 1=red, 2=yellow, and 3=blue (score error). The *$scorecolor* variable will not show a number for the gray (disabled) status, because an agent with a gray status is inactive, and will not return action file variables.
$redThreshold	The value at which the status of the agent turns red
$greenThreshold	The value at which the status of the agent turns green

These variables can be included in an action file as part of an email message to provide descriptive information about the threshold crossing.

Examples

Suppose you are creating an action file for a Partition Utilization agent that monitors all partitions in the traffic tree and you want to be notified via email when the utilization of one or more partitions exceeds the defined utilization threshold. The action file would contain the following command:

```
send email bob@test.com "$agentname agent turned red" "$scorevalue classes crossed the Utilization threshold."
```

When the agent turns red and the action file gets executed, an email will be sent to bob@test.com. The subject of the email will be *Partition Utilization agent turned red* and the body of the message will have something like *2 classes crossed the Utilization threshold*.

Agent-Specific Variables

Any parameter for an agent can be a variable for an action file. For example, the Failed Flow Ratio agent has parameters named RedThreshold, GreenThreshold, RatioThreshold, ViolatingHosts, and ExceptionHosts. To use these as variables for action files, add a dollar sign in front of the parameter name ($RedThreshold, $GreenThreshold, $RatioThreshold, $ViolatingHosts, $ExceptionHosts). In addition to these parameter-based variables, some agent types offer additional variables that can be included in action files.

Agent Template	Parameters
Class ME Variables Link ME Variables Partition ME Variables	$className: The value of the ClassName parameter. ClassName is the name of the traffic class the agent is monitoring. $meVariableName: The value of the MeVariableName parameter. MeVariableName is the name of the measurement variable the agent is measuring. $meDuration: The value of the MeDuration parameter. MeDuration is the period of time over which measurements will be taken, in minutes. $greenOperator: The value of the GreenOperator parameter (for example, LT for less than, EQ for equal to, GTE for greater than or equal to) $redOperator: The value of the RedOperator parameter $class-id: The name of the traffic class that the agent is monitoring $me-value: The value of the measurement variable being measured by the agent, as recorded at the end of the evaluation interval $me-variable: The name of the measurement variable being measured by the agent (Note: If the measurement variable is derived from calculations involving other variables, the name of the variable will not be displayed. Instead it will say *Derived*.)
Default Traffic	$className: The value of the ClassName parameter. ClassName is the name of the traffic class the agent is monitoring. $avg-bps: Usage of the class being monitored, averaged over the evaluation interval, measured in bits per second $class-id: The name of the traffic class the agent is monitoring
Failed Flow Ratio	$ratioThreshold: The value of the RatioThreshold parameter. RatioThreshold is the percentage of new client flows per minute that are failed flows; when a host crosses this threshold, it is considered to be in violation $violatingHosts: The name of the host list that contains the IP addresses of the hosts that were in violation (including hosts from previous intervals unless the hosts have been removed) $exceptionHosts: The name of the host list that contains the IP addresses of the hosts in the host exception list

Agent Template	Parameters
High Bandwidth Host	$host-ip: The IP address of the top usage host, in dotted-decimal format. This variable is valid in red action files only. $direction: The direction of the top usage host's traffic (/Inbound or /Outbound). This variable is valid in red action files only. $avg-bps: Usage of the top usage host, averaged over the evaluation interval, measured in bits per second. This variable is valid in red action files only. $violatingHosts: The name of the host list that contains the IP addresses of the hosts that were in violation (including hosts from previous intervals unless the hosts have been removed) $exceptionHosts: The name of the host list that contains the IP addresses of the hosts in the host exception list
High Bandwidth New App	$namelist: The name of the traffic class with the highest bandwidth usage that was auto-discovered during the evaluation interval $avg-bps: Usage of the $namelist class, averaged over the evaluation interval, measured in bits per second
Host Info Variables	$variableName: The value of the VariableName parameter (Current Connections, New Flows Client, New Flows Server, Failed Flows) $flowsThreshold: The value of the FlowsThreshold parameter. FlowsThreshold is the number of new flows or connections per minute; when a host crosses this threshold, it is considered to be in violation. $violatingHosts: The name of the host list that contains the IP addresses of the hosts that were in violation (including hosts from previous intervals unless the hosts have been removed) $exceptionHosts: The name of the host list that contains the IP addresses of the hosts in the host exception list
NFPM Failed Flow	$flowsThreshold: The user-defined value for the FlowsThreshold parameter. FlowsThreshold is the number of new flows per minute; when a host crosses this threshold, it is considered to be in violation. $violatingHosts: The name of the host list that contains the IP addresses of the hosts that were in violation (including hosts from previous intervals unless the hosts have been removed) $exceptionHosts: The name of the host list that contains the IP addresses of the hosts in the host exception list
NFPM Side Unknown	$side: The value of the Side parameter (one of the following: client, server) $sideThreshold: The value of the SideThreshold parameter. SideThreshold is the number of new flows per minute; when a host crosses this threshold, it is considered to be in violation. $violatingHosts: The name of the host list that contains the IP addresses of the hosts that were in violation (including hosts from previous intervals unless the hosts have been removed) $exceptionHosts: The name of the host list that contains the IP addresses of the hosts in the host exception list

Agent Template	Parameters
Partition Utilization	$className: The value of the ClassName parameter. ClassName is the name of the traffic class the agent is monitoring. $utilization: The value of the Utilization parameter. Utilization is the percent utilization of the partition; when a partition exceeds the defined Utilization, the class is considered to be in violation. $avg-bps: Usage in the violating class' partition, averaged over the evaluation interval, measured in bits per second. This variable is valid in red action files only. $class-id: The name of the traffic class (partition) that caused the agent to cross the red threshold. This variable is valid in red action files only. **Note:** The $avg-bps and $class-id variables work well in action files when the agent is monitoring a single partition. These variables are not intended for use when the agent is monitoring multiple classes (that is, when ClassName contains a wildcard).
Quota Bandwidth Host	$bytes: The number of bytes a violating host consumed during the monitored interval $host-ip: The IP address of the violating host, in dotted-decimal format $direction: The direction of the violating host's traffic $violatingHosts: The name of the host list that contains the IP addresses of the hosts that were in violation (including hosts from previous intervals unless the hosts have been removed) $exceptionHosts: The name of the host list that contains the IP addresses of the hosts in the host exception list
Traffic Performance	$className: The value of the ClassName parameter. ClassName is the name of the traffic class the agent is monitoring. $efficiency: The value of the Efficiency parameter. Efficiency is the transmission effiency of the traffic class; when a class drops below the defined Efficiency percentage, it is considered to be in violation. $network-efficiency: The percentage of bytes in the violating class that were *not* retransmitted $class-id: The name of the traffic class that caused the agent to cross the red threshold **Note:** The $network-efficiency and $class-id variables work well in action files when the agent is monitoring a single class. These variables are not intended for use when the agent is monitoring multiple classes (that is, when ClassName contains a wildcard).

Examples

An action file for the High Bandwidth Host agent could contain the following command, which makes use of two of its agent-specific variables ($direction and $host-ip):

```
class new $direction $host-ip outside host:$host-ip
```

This action file creates a class for the top host in the interval.

As another example, suppose you have created a Partition Utilization agent to monitor the /Inbound/FTP partition, and you entered the traffic class */Inbound/FTP* as the agent's *ClassName* parameter when you defined the agent. The action file for this agent could include the following CLI command:

```
partition apply $className 30% 40%
```

When the action file runs, the name of the class specified in the agent's ClassName parameter will automatically be inserted into the action file. Thus, the action file would then execute the CLI command:

```
partition apply /inbound/FTP 30% 40%
```

If you ever change the ClassName parameter on the agent to measure a different partition, for example on the */Inbound/Citrix* traffic class, the action file would automatically reflect the change, and when triggered, would execute the command:

```
partition apply /inbound/Citrix 30% 40%
```

User-Defined Variables

The adaptive response feature also allows you to specify your own unique action file variables. When you define your own variables in an action file, these variables will appear in the *Edit Agent Entry* window along with the standard parameters (RedThreshold, GreenThreshold, and so forth). **show screen...**

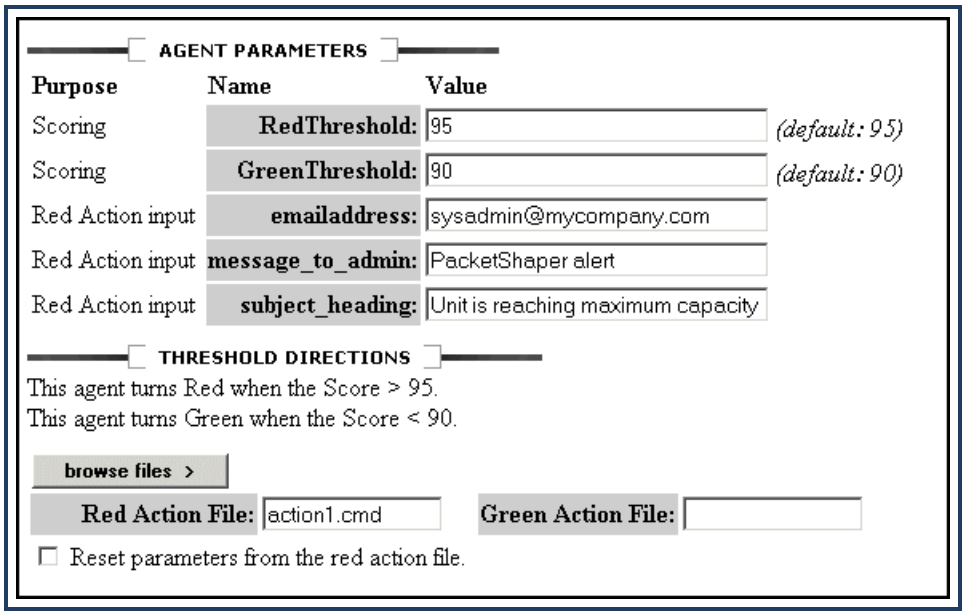

The advantage to using user-defined variables in an action file is that if information changes, you don't need to modify the action file you can just directly enter the changed values in the fields in the *Edit Agent Entry* window.

A user-defined variable:

- Must begin with a dollar sign ($)
- Can be up to 32 characters long (not including the $ sign)
- Can have alphanumeric characters

- Can contain the following special characters: underscore (_), hyphen (-), percent (%)
- Cannot contain spaces

Examples of valid variable names: $email_address, $email2, $partition-maximum%.

Example

A red action file could include the following CLI command:

```
send email $emailaddress $subject_heading $message_to_admin
```

As a result, these user-defined action file variables will appear on the *Edit Agent Entry* window with a corresponding field that lets you enter a value.

> **Note:** Before the user-defined variables appear on the *Edit Agent Entry* screen, you must reset the parameters.

Adaptive Response Agents Overview

Adaptive response agents monitor your PacketShaper unit, network health, hosts, and applications, returning values and evaluating them against their predefined green and red performance thresholds. If there is a problem, the agent will send information about the situation to two different places: an incident report and an action file.

PacketShaper has a set of agents that are preconfigured, or you can define your own agents based on a variety of templates. Click any of the adaptive response agent template names below for additional information on that template, including default thresholds, parameters, and action file variables.

Unit Health Agent Templates

- Unit Limits
- System Load
- Memory Allocation

Network Health Agent Templates

- Traffic Performance
- Partition Utilization

Application Health Agent Templates

- Default Traffic
- New Application
- High Bandwidth New App

Host Agent Templates

- High Bandwidth Host
- New Flows Per Minute (NFPM) Side Unknown
- New Flows Per Minute (NFPM) Failed Flow
- Failed Flow Ratio
- Host Info Variables
- Quota Bandwidth Host

User Event Emulation Agent Templates

- Class ME Variables
- Partition ME Variables
- Link ME Variables

BCAAA Overview

PacketShaper uses the Symantec Authentication and Authorization Agent (BCAAA) to resolve IP addresses to user names so that it can classify and report on network users and user groups. BCAAA is a separate software package supplied to Symantec customers who have a valid support contract.

BCAAA is installed on a domain member server and uses Windows APIs to communicate with the Microsoft Active Directory (AD) Domain Server. The BCAAA application, acting as a secure intermediary service between the PacketShaper and AD, passes user identification from AD to PacketShaper, allowing PacketShaper to identify traffic flows by users and groups. After BCAAA is installed and configured, PacketShaper administrators can create user-based reports to monitor usage and define policies to manage access to applications, both internal and external. The ability to employ user identification information in the reporting and policy making process allows administrators to have more granular visibility and control over connections and content.

How Does PacketShaper Work with BCAAA?

PacketShaper uses BCAAA to retrieve the user identification from the Active Directory Domain Server(s) for IP addresses generating requests to application and web servers. By correlating user identification with IP addresses, you can apply PacketShaper policies to WAN and Internet traffic by user and group.

Upon seeing a new traffic flow from an IP address, PacketShaper initiates a request to BCAAA to get a user name and the names of all groups the user belongs to. BCAAA, via Windows APIs, queries the Active Directory Domain Server, obtains the user identity, and provides it to PacketShaper. PacketShaper then updates its identification of the network traffic flow from IP address to user name and group names. If a user drops off the network and another user takes over the IP address via dynamic allocation, then the PacketShaper query to BCAAA produces a new user ID for the IP address. If a user creates multiple traffic flows from the same IP address, each flow is separately identified by User ID.

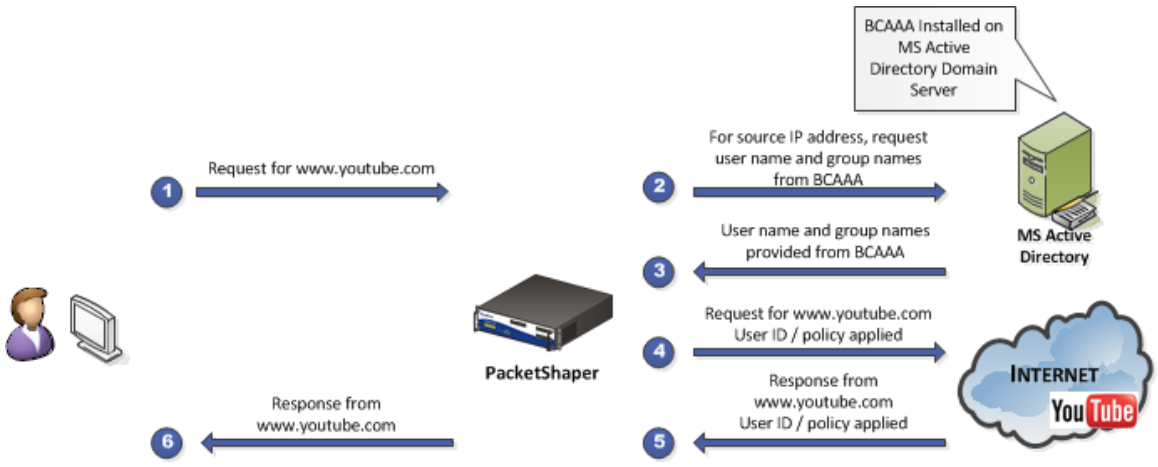

User Caches

As PacketShaper identifies traffic, it queries BCAAA to see if the source or destination IP address is associated with a user name. PacketShaper stores the query results in two local caches. When the query results in an IP-user mapping, this information is stored in the IP-user cache. If the IP address was looked up in BCAAA and there was no user name associated with the IP address, this information is stored in an IP-only cache.

By caching the IP addresses and user and group names, PacketShaper is able to speed up user name lookups, as querying the on-box cache is faster than querying the external BCAAA server. When classifying traffic, PacketShaper will look for the IP address in both caches. If the IP address has previously been associated with a user and is in the IP-user cache, it will identify this traffic by user name. If the IP address is in the IP-only cache, PacketShaper knows that there isn't a user name associated with the IP, so it doesn't waste time querying BCAAA. If the IP address is not in either cache, PacketShaper will query BCAAA for the user name association.

The IP addresses that have a user mapping are stored in the cache for one hour by default; IPs with no user mapping are stored for 30 minutes. You can adjust the cache timeout settings on theSystem Variable setup page. Because querying the cache is faster than querying the BCAAA server, you can accelerate user name lookups by increasing the cache timeout. However, the tradeoff is that stale mappings could cause incorrect user name identification. For relatively static network environments, storing user mappings for one hour should be sufficient; for more dynamic environments (BYOD, laptops), you may want to store user mappings for a much shorter time period, such as five minutes.

Planning for your BCAAA Installation

It's important to size your BCAAA installation appropriately when preparing to use the user awareness feature. To ensure that BCAAA can service all the queries from PacketShaper, Symantec recommends the following BCAAA-to-PacketShaper ratios:

- Install one instance of BCAAA per PacketShaper, for any number of supported hosts.

  or

- For redundancy, a PacketShaper can connect to two BCAAA servers. If PacketShaper detects that the primary BCAAA server is down, it automatically connects to the secondary BCAAA server. One server is active, and the other is standby.

or

- Multiple PacketShapers can share a single BCAAA server if the number of active hosts on each PacketShaper is less than or equal to the IP-user cache size for that model (equivalent to three times the supported number of users).

Disregarding these recommendations could overload the BCAAA server, resulting in BCAAA being unable to respond to PacketShaper IP-user queries. See Configuration Limits in PacketGuide for per-model specifications.

Caution: To prevent any impact on ProxySG authentication functionality, PacketShaper should not share a BCAAA server with a ProxySG.

Bypass Mode Overview

Bypass mode is an operational state in which the PacketShaper connects the Inside and Outside Ethernet ports directly to each other, essentially behaving like a piece of cable. By default, the appliance enters bypass mode when PacketShaper loses power for some reason, for example, due to a power outage or unit failure.

With PacketShaper's default bypass settings, if an appliance loses power, relays close on a NIC to allow user traffic to pass unhindered between the two interfaces.The setting for this behavior is *bypass closed* and the appliance is engaged in *bypass mode* if the appliance loses power. If you want PacketShaper to be a point of failure in the network when the appliance is powered off or goes down, you can set bypass to open (see setup bypass).

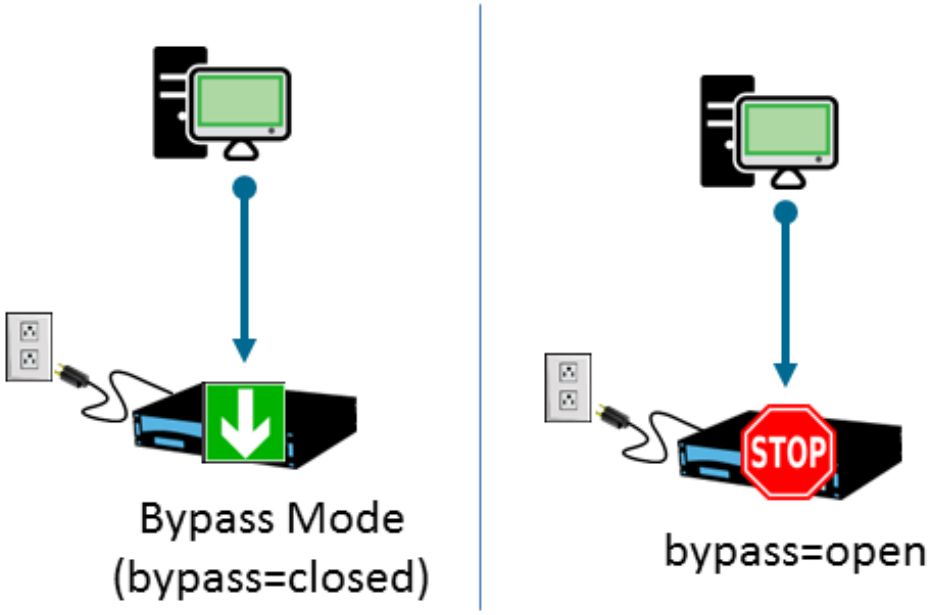

Traffic will only be blocked for a few seconds while the appliance is resetting or rebooting.

Captive Portal Overview

PacketShaper's *captive portal* feature allows enterprises and service providers to present an Acceptable Use Policy (AUP) to guest users accessing the Internet in a web browser on a WiFi network. Users with source IPs in the defined restricted subnets will see an AUP message they must accept before they are granted Internet access. If the user does not agree to the terms and conditions of the AUP, the host will be blocked on the network. The AUP page is prompted every time a blocked host attempts to connect to a website in a browser. A typical use of the captive portal feature is the guest WiFi network in a hotel or airport, where users are required to agree to the terms and conditions before they can use the network.

1. User with a source IP in the restricted subnet enters an HTTP URL, such as cnn.com, in a web browser.

   Note: If the user enters an HTTPS URL, the connection will be dropped.

2. PacketShaper redirects to the AUP message and displays the page in the user's web browser.

3. User clicks **Agree** to accept the terms and conditions.

4. PacketShaper redirects to symantec.com and displays this website in the user's browser.

5. User reenters the original URL and now has Internet access on the WiFi network.

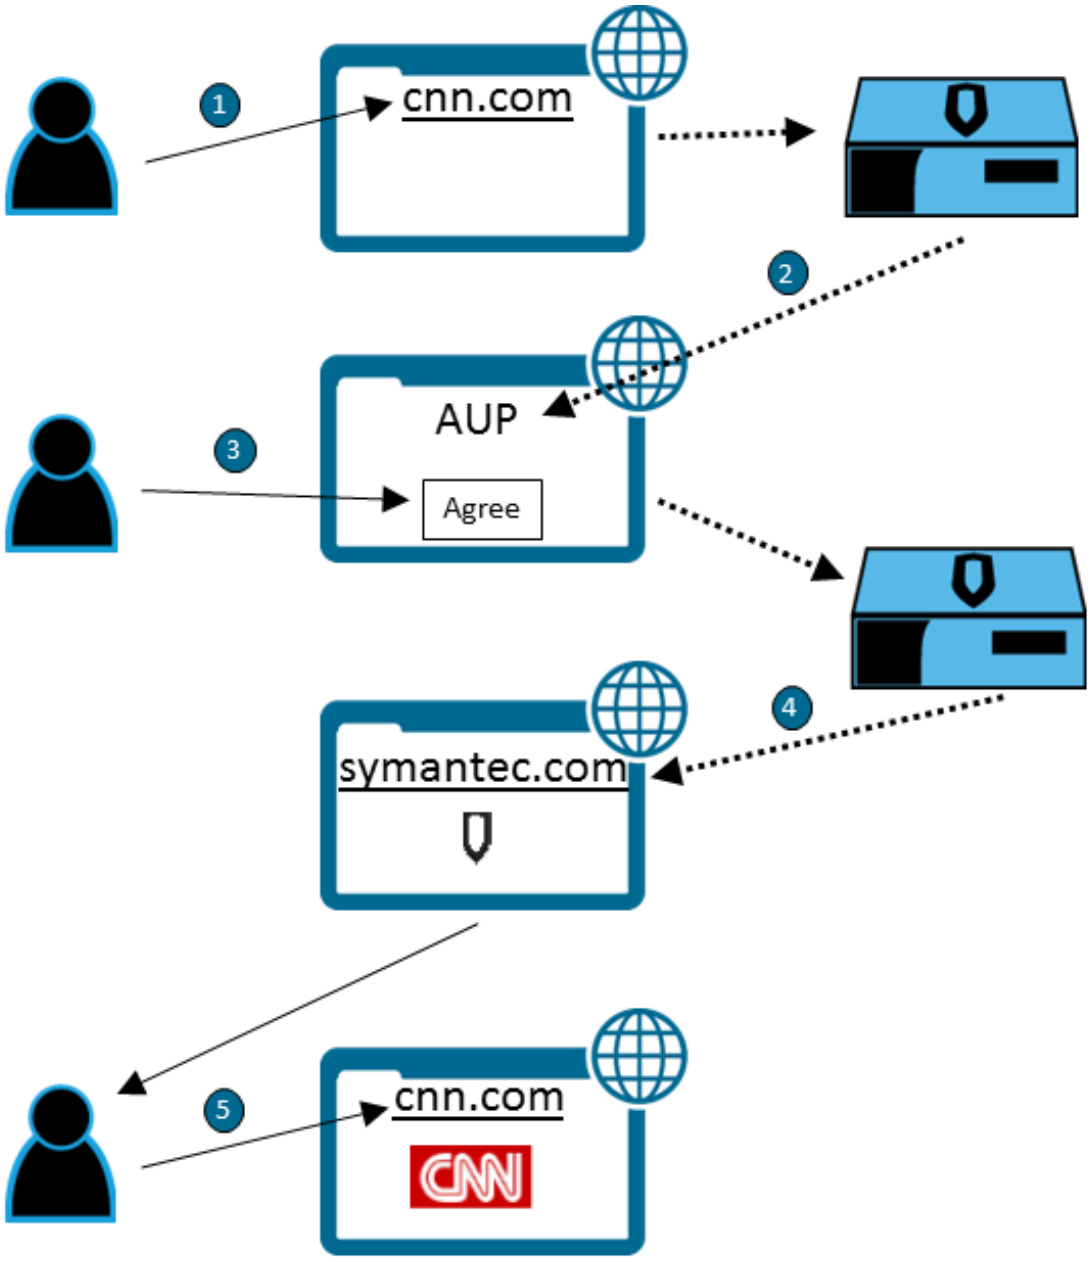

After a user agrees to the AUP, the session is good until midnight of the first day of the next calendar month, based on the PacketShaper's local time.

If an enterprise has multiple WiFi subnets, the administrator can create a captive portal *site* for each network segment; each site can optionally have a custom AUP. In this type of captive portal implementation, these sites are all associated with the same *customer* ID. In the case of a service provider using the captive portal feature, there would be multiple customers, and each customer could have multiple sites. Each customer has a default AUP that applies to all sites, unless a custom AUP is specified for the site; in other words, site-specific AUPs override the customer AUP.

Captive Portal Feature Requirements

- Because captive portal AUP is triggered by HTTP redirection and rendered through HTTP, HTTP access on PacketShaper must be enabled. Use the **`setup access enable HTTP`** command in the CLI.
- To share the captive portal settings with other PacketShapers, the PolicyCenter appliance must be running a version that supports the captive portal feature. Shareable settings include the enable/disable status, redirection IP addresses, and the starting IP address and subnet mask of the guest WiFi subnet. Customer and site configuration is specific to a local PacketShaper and is not shareable.

Additional Information

- The captive portal AUP page is displayed on the first HTTP connection from the managed host. No HTTPS redirection is available.
- All captive portal settings are persistent through appliance reboot. However, AUP acceptance status is not persistent. AUP will be reprompted after PacketShaper reboot regardless of past acceptance.

Traffic Classification Overview

Terms

For traffic classification, it is helpful to understand these key terms.

Term	Definition
Class	A logical grouping of traffic flows that share the same characteristics and are analyzed or controlled together.
Matching Rule	The portion of a class' definition that specifies the type of traffic associated with the class a specific application, protocol, address, or set of addresses. See Matching Rules.
Partition	A bandwidth pipe assigned to a given class to protect or restrict all the flows in that class. See Partition Overview.
Policy	A rule assigned to a given class that defines how a single flow will be handled during bandwidth management. Policies act on individual traffic flows, for example an HTTP connection to a web server. See Policy Overview.

Why Classify Traffic?

In order to analyze or control one type of traffic distinctly, PacketShaper must be able to differentiate it from other types of traffic. The traffic tree's traffic classes are PacketShaper's mechanism for identifying and organizing different types of traffic.

For example, if you want to measure bandwidth utilization for SAP and Oracle, then your traffic tree must distinguish SAP from Oracle traffic and have separate classes for each. Or, if you want to allocate set amounts of bandwidth to your Paris and Oslo offices, then your tree must separate Paris traffic from Oslo traffic and have separate classes for each.

PacketShaper differentiates one application from another by evaluating characteristics in traffic flows and organizing them into *classes*. Each traffic class contains at least one *matching rule*, a set of characteristics that identifies a specific traffic type.

While many applications or devices can identify traffic on well-known ports, PacketShaper is "application aware" and goes beyond classification by port number. For example, QuickTime can be embedded within HTTP and generally could be categorized as web traffic on port 80; however, PacketShaper is able to differentiate QuickTime traffic from other web traffic.

Creating Traffic Classes

Traffic classes can be created automatically, using the traffic discovery feature, or manually.

- To enable the traffic discovery feature, see Enable/Disable Traffic Discovery.
- To create a traffic class manually, see Create a Traffic Class.

Control Strategy Overview

Controlling bandwidth allocation for today's traffic diversity is a combination of an art and a learned skill. Symantec recommends considering the following sequence and points when approaching control features for the first time.

These points apply to situations where you are managing the traffic and applications to and from one location. If you are managing multiple locations with a single PacketShaper, Control Branch Traffic from a Main Site might be more useful.

1. First, develop a useful traffic tree that has traffic classes for the types of traffic you want to control.

2. Determine which applications are undermining the performance of others.

3. Identify the applications whose performance is crucial to your organization.

4. Use the *Monitor Traffic* window to sort the list by peak bps. Make a list of any class that has peaked to more than 50 percent of the link size.

5. Using your three lists of applications those impacting others' performance, those that are important, and those that peak at more than 50 percent of link size characterize your traffic, categorizing on the basis of importance, size, and sensitivity to latency and jitter.

   It's not necessary to assign a policy to every traffic class. (It's even advisable to avoid over-managing your traffic by creating too many policies and partitions on too many types of traffic.) If you don't assign a specific policy to a class, the class inherits the Inbound or Outbound Default class policy, which is set at the factory as a priority policy with a priority of 3. Although you can change the policy for a Default class, it is *not* recommended, as this can cause unpredictable results.

   Start by managing only your mission-critical traffic and the traffic that impacts others' performance.

6. Define partitions and/or policies to protect your important traffic classes. As resources, you can use:

   - Policy/Partition Guidelines (tables of policies, partitions, and guidelines about when they would and would not be appropriate)
   - Suggested Policies (on-line, in-product policy suggestions)
   - Reserve bandwidth (instructions for protecting bandwidth for the total of all traffic in an application)
   - Insulate users of an application from their peers (instructions to control each application session)

7. Define partitions and/or policies to contain your important traffic that impacts others or bursts to a large portion of the link. As resources, you can use:

   - Policy/Partition Guidelines (tables of policies, partitions, and guidelines about when they would and would not be appropriate)

- Suggested Policies (on-line, in-product policy suggestions)
- Contain Greedy Business Applications (instructions for limiting an important application)

8. Define partitions and/or policies to contain your unimportant traffic that impacts others or bursts to a large portion of the link. As resources, you can use:

   - Policy/Partition Guidelines (tables of policies, partitions, and guidelines about when they would and would not be appropriate)
   - Suggested Policies (on-line, in-product policy suggestions)
   - Limit Total Bandwidth (generic instructions for limiting with a partition)
   - Block Unwanted Traffic (instructions for completely stopping one type of traffic)

9. When in doubt, follow these rules:

   - Don't over manage. Reign in the top pest applications. Protect the top critical applications. And stop.
   - If you care about your end users' experience and the traffic has large and long flows, use a rate policy with a minimum guaranteed rate of 0 and no guaranteed rate limit, and assign an appropriate priority for bursting.
   - If you care about your end users' experience and the traffic has small and short flows, use a priority policy with an appropriate priority.
   - If you do not care about your end users' experience, use a partition with a limit.

Other Useful Considerations

Most protocols are asymmetric, so a policy may fit a profile for one direction of the traffic flow, but a different policy may be required for control of the traffic in the opposite direction.

In general, there are four reasons to control an application:

1. **To provision** bandwidth for applications that need a set per-session amount to perform (VoIP, Citrix)
2. **To protect** mission-critical applications
3. **To contain** bandwidth-greedy applications that impact others
4. **To balance** necessary and important (but bandwidth-greedy) applications such as HTTP.

In certain cases, the location of the server may govern how you manage the traffic. For example, if you are hosting a web server at your site, you may want to give this traffic a higher priority than your users' web-browsing sessions. You can create a customized HTTP traffic class that references the location of this inside server, then use a policy to give this traffic priority. For details about inside and outside servers, see Server Location.

Remember that a policy controls how much bandwidth is allocated to each traffic flow. Do *not* assign large amounts of guaranteed bandwidth to a traffic class because, as the number of simultaneous sessions increases, you'll run out of bandwidth. To determine if you have exhausted available bandwidth for a class, check the guaranteed rate failures in the *Monitor Traffic* window. For details, see Monitor Traffic.

Customer Portal Overview

The customer portal feature offers service providers the ability to design custom network and application status pages that can be viewed by their customers. Customer portals allow for providers to offer customized, limited, branded sets of PacketShaper information to their end users—information such as throughput, response times, and other metrics collected by PacketShaper, as well as statistics that indicate whether Service Level Agreements (SLAs) are being followed.

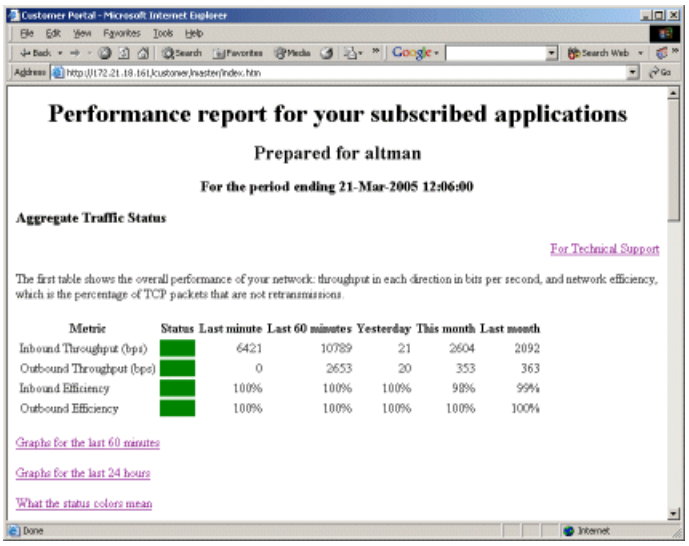

Unique user accounts and directories are created for each customer account. The service provider can either create a single HTML page that has customized data for each portal user, or create an individual HTML page for each customer. The maximum number of portal pages you can create is limited by the maximum number of classes your unit supports. For example, PS-S500's maximum number of classes is 10,000, so you can create up to 10,000 different portal pages.

Follow these steps to configure and use the customer portal feature:

1. Create user accounts, and optionally configure individual or global message-of-the-day settings. See " Create Customer Portal Accounts" on page 748.

2. Assign an IP address to the customer portal. See "Assign a Customer Portal Address" on page 758.

3. Download sample files from Symantec to use as templates in designing your own login and portal pages. See " Download Sample Portal Files" on page 754.

4. Customize the customer login screen. See " Customize the Login Screen" on page 756.

5. Design custom HTML pages to display the desired statistics for the users of the customer portal. See " Design Customer Portal Pages" on page 757.

6. Log into the customer portal. See " Log Into the Customer Portal" on page 759.

If you want customers to use RADIUS authentication when logging into the portal, you need to configure a RADIUS authentication server.

Overview of Event Notification

A PacketShaper *event* is a network situation that can trigger an alert message sent via an email, SNMP trap, Syslog server, or a combination of any of these methods. The event is based on a particular measurement variable, associated with a specific class, partition, or link, and has a defined threshold. When the threshold value is exceeded, PacketShaper will automatically send notification of the event and record the event in a log. For example, you might have an event defined to notify you when the percentage of TCP connection aborts (the tcp-conn-aborts% variable) exceeds 70% (the threshold value). In this example, the event can notify you of conditions that could lead to poor customer service.

Network and system administrators rely on event notification to proactively manage network resources. PacketShaper's event notification feature takes advantage of the network metrics already collected by the unit. For example, you can use event notification to automatically report to you when there are unacceptable network delays or when there isn't enough bandwidth to satisfy guaranteed rate demand. By having PacketShaper send you notifications automatically, you have the opportunity to be aware of situations before they trigger user complaints.

Using PacketShaper's event notification feature, you can:

- Define your own events or use one of the predefined events.
- Configure thresholds for key performance variables, which are sampled at regular intervals.
- Receive notification emails, SNMP traps, and/or Syslog messages when threshold crossings occur. Notifications will also be sent out when the situation returns to normal.
- Access event log files stored on the unit's data disk.

Configuring the Event Feature

Before creating events, you need to decide how you want to be notified when events occur. To be notified via email, you must identify the email server and email recipients. If you want PacketShaper to send an SNMP trap, you must define the SNMP trap destination.

Setting Up and Managing Events

Once you have finished setting up event notification, you can start configuring events. Event notification is configured via the command-line interface (CLI); after events are configured, event status and statistics can be monitored through the browser user interface.

> **Note:** The User Event facility may not be accessible for a few minutes after a measurement engine reset.

Flow Detail Records Overview

The Flow Detail Records (FDR) feature is a method for gathering and processing per-flow statistics. When FDR is enabled, the PacketShaper will become an *emitter*, periodically pushing data to a remote system called a *collector*. The unit will emit records that contain details of all flows that go through PacketShaper to a collector, such as Symantec s IntelligenceCenter. These records are called *flow detail records*.

PacketShaper can look at a flow, identify its application or protocol, gather statistics about this flow, include this information in the PacketShaper flow detail record, and then send the FDRs to a collector. In the collector s report generator, you can view reports to summarize and analyze the data.

Symantec s FDR solution offers:

- **Enhanced troubleshooting and forensic capabilities.** IntelligenceCenter Reports can aid in troubleshooting network problems and help determine the source of a DoS attack. For example, the IntelligenceCenter host reports let you see the busiest hosts on your network and drill down to see the applications used and destination addresses contacted.

- **Integration with accounting/billing programs.** Broadband service providers can bill customers by application usage and, if desired, have different billing rates for different types of applications (such as P2P, VoIP, email, and web surfing). Or, enterprises can track each department s application usage and bill them accordingly.

- **Statistics of latency, loss, and utilization of VoIP flows.** You can use these statistics to monitor and analyze the effectiveness of VoIP on your network.
- **Historical Top Talker and Top Listener data.** In IntelligenceCenter, you can specify how long you want to keep the data; for example, you may want to keep hourly data for six weeks and monthly data for two years. You also have an increased number of Top Talkers/Listeners (PacketShaper offers a maximum of 20 hosts per traffic class; in IntelligenceCenter, this value is user specified.)
- **Cisco NetFlow v5 compatibility.** PacketShaper can emit FDRs to NetFlow-compatible collectors and analyzers, such as Evident Billing Software and Cisco Collector.

What Type of Information is in a Flow Detail Record?

Generally, a flow detail record (FDR) contains information about a TCP or non-TCP flow, such as source and destination IP addresses, the size of the flow (in terms of packets and bytes), and when the flow was sent. The specific fields of information vary according to the type of record format. PacketShaper offers three different record types: Packeteer-1, Packeteer-2, and NetFlow-5. The NetFlow-5 record type identifies the flow s Layer 4 protocol (such as TCP, UDP, or ICMP) and IP ToS/Diffserv. The Packeteer-2 format contains all the NetFlow fields as well as PacketShaper-specific data, for example: the traffic class into which the flow was classified, type of policy, number of retransmitted bytes, Response Time Measurement (RTM) data, packet exchange time, and VoIP statistics for RTCP VoIP streams. The Packeteer-1 format can be emitted, but collectors for this format are not currently available.

For a list and description of all fields contained in in the NetFlow-5 record type, see NetFlow v5 Record Format.

Flow Detail Record Collectors

A flow detail record collector is a software application, such as IntelligenceCenter, that accumulates the data from an FDR emitter (PacketShaper). Most collectors do much more than gather the data they also massage and present the information in a meaningful way in reports and graphs.

Use the Packeteer-2 format to send flow detail records to the IntelligenceCenter collector. Use the NetFlow-5 record type to send flow detail records to a NetFlow v5 collector.

In the example below, IntelligenceCenter has been set up as a collector that uses the Packeteer-2 format, and Cisco Collector and Evident Billing Software have been defined as collectors of the NetFlow-5 format.

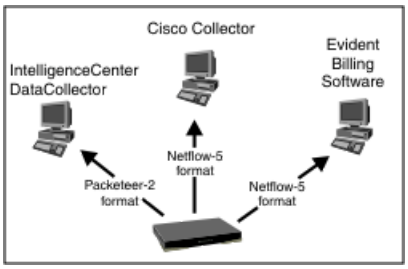

How Often are Flow Detail Records Emitted?

PacketShaper can emit flow detail records either at the end of a flow (default) or at a set interval (optional):

- End-of-flow FDRs
- Intermediate FDRs

> **Note:** Changes made to FDR behavior impact the records emitted to all FDR collectors. For example, you cannot configure your unit to emit intermediate FDRs only to a NetFlow-5 collector, while sending end-of-flow FDRs to a Packeteer-2 collector, such as IntelligenceCenter (which cannot process intermediate FDRs).

End-of-flow FDRs

For TCP flows, two flow records (one for the start of the flow, one for the end) are sent when the TCP connection is closed. In the unusual case when connections remain open for a long period of time without any activity, PacketShaper will eventually reclaim the resources and close the connection; the flow records will be created at that time.

For non-TCP flows, flow records are generally created one hour after PacketShaper sees the last packet for the flow. Exceptions are transactional non-TCP flows, such as a DNS lookup over UDP or an ICMP ping. For these types of flows, the flow record is created when the transaction is completed.

Flow detail records are bundled into UDP packets before they are emitted to the collector. Because of this bundling process, there is a short delay from the time flow detail records are created until the UDP packet is emitted. On a busy PacketShaper, this delay is typically less than a second.

UDP Flow Record Packet for NetFlow v5

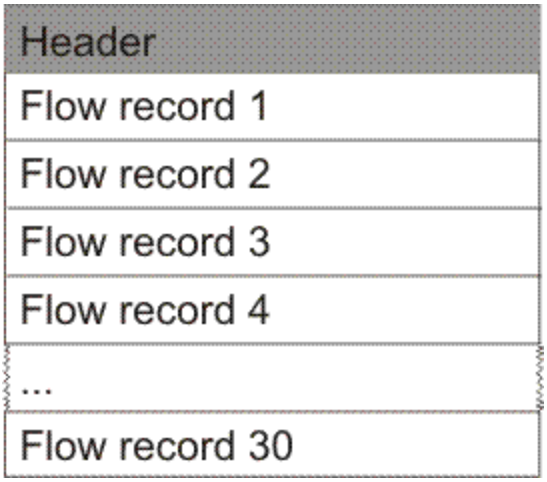

> **Note:** If flow recording is on and then turned off, any unsent flow records are bundled into a UDP packet and emitted to the collector as recording is turned off.

At the top of each UDP flow record packet is a header. The header provides information about the records in the UDP packet, such as the number of flow records in the packet. For a list and description of all fields contained in the NetFlow header, see UDP header for NetFlow v5 packets.

Intermediate FDRs

When the intermediate FDR feature is enabled, FDRs are emitted at a configurable interval throughout long TCP flows as well as at the end of flows. This allows a suitably-instrumented collector, such as Cisco-based Netflow-5 collectors, to report flow data during long-lived flows. For information on how to enable intermediate FDRs, see Enable Intermediate FDRs.

> **Note:** IntelligenceCenter does not support intermediate FDRs.

Classification of Flow Detail Records

Flow detail record packets are automatically classified into two predefined classes:

- **FlowRecords** – includes Packeteer-1 and Packeteer-2 FDRs
- **NetFlowV5** – includes NetFlow v5 FDRs emitted by the PacketShaper

The appropriate class is automatically created as soon as a collector is enabled. The FDR-related classes are child classes of Outbound/Localhost because FDRs are generated by PacketShaper (the local host). For example, if a Packeteer-2 collector is defined and enabled, an Outbound/Localhost/FlowRecords class is created; this class tracks all FDR traffic emitted to the defined Packeteer-2 collector.

The service names correspond to the class names: FlowRecords and NetFlowV5. If you delete the classes that were automatically created, you can either manually recreate them, specifying FlowRecords or NetFlowV5 for the service, or reset the unit (the classes will be automatically recreated after a reset).

Feature Requirements and Limitations

The FDR feature has several differences from Cisco's NetFlow:

- PacketShaper does not support aggregation (for example, aggregate all flows to the same destination IP and report them as one flow) or sampling (collect details on every nth packet).
- For short-lived TCP flows, PacketShaper reports the flow data as soon as the connection is closed. NetFlow typically reports it about 15 seconds later.
- For short UDP flows (such as a 30-second Voice over IP conversation), PacketShaper closes the flow and reports it after one hour (when PacketShaper times out non-transactional UDP flows). NetFlow reports the flow in about 15 seconds.

IntelligenceCenter Overview

IntelligenceCenter (IC) is a central management solution for Symantec products. IC provides a secure dashboard that allows access to Symantec products from a central screen. With a single click, you can go into PacketShaper—without having to enter a user name and password each time a product is accessed. Just log in to IC and the Symantec products are instantly accessible without additional sign on.

DataCollector (DC) is another component of IntelligenceCenter. You can configure DC to collect Measurement Engine (ME) and/or Flow Detail Record (FDR) data from PacketShaper and iShaper appliances on your network. You can also collect NetFlow-5 data from other network devices such as routers. You can then use the data to generate aggregated reports that provide snapshots of your network as a whole. Alternatively, you can define your own applications and network views, which allow you to generate reports using specific cross-sections of data.

IC provides two types of reporting mechanisms:

- Static Reports—Allow you to choose from a selection of predefined reports, which you can run ad hoc or on a schedule.
- Portlets—Allow you to monitor and troubleshoot the flows and applications that are the most important to you. Portlet configurations are unique to each IC user profile.

IntelligenceCenter provides a comprehensive set of reports that allows you monitor the performance of the applications, devices, and hosts that are running on your network. The data that IntelligenceCenter displays is collected from the PacketShaper appliances throughout your network and aggregated according to options you define when you run a report. You can generate reports that include data for your entire network, or you can define a specific geographic region or logical group – such as a business unit – that is of interest to you. You can also generate reports for a specific site (subnet class).

IPv6 Classification Overview

PacketShaper can identify IPv6 traffic and sub-classify a variety of IPv6 protocols, such as CIFS and HTTP. IPv6 traffic is classified into separate services from IPv4 traffic, for example HTTP-IPv6, FTP-IPv6, and CIFS-IPv6.

When traffic discovery is enabled, IPv6 classes are auto-discovered in the Inbound and Outbound branches of the traffic tree just as IPv4 classes are. You can also manually create IPv6 classes in the Advanced UI, the Sky UI, and the CLI.

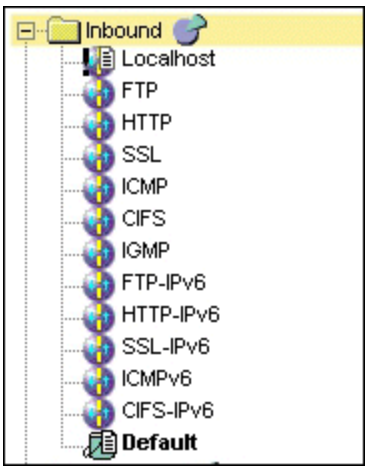

The tree above shows IPv4 and IPv6 auto-discovered or manually created at the root. However, there are other ways to build your traffic tree around IPv6 classification. Refer to the following topics:

- Create a Site-Based Traffic Tree for IPv6 Subnets
- Create a Service-Centric Traffic Tree for IPv6
- Create an IPv6-Based Branch in the Traffic Tree

Control of IPv6 Traffic

Controlling IPv6 traffic is no different than controlling IPv4 traffic. First, you define classes for the type of traffic you want to control. Next, you assign appropriate policies and partitions; you can look at PacketShaper reports to see current utilization levels of various types of IPv4 and IPv6 traffic. Finally, you make sure shaping is enabled. You can then review your reports to see how shaping has affected the rate and distribution of traffic on your network.

Features that Support IPv6

The following PacketShaper features support IPv6:

- Auto-discovery of IPv6 traffic
- Manual creation of IPv6 classes (IPv6 services, addresses, subnets) in Advanced UI, Sky UI, and CLI

- Setting of partitions, dynamic sub-partitions, and policies on IPv6 classes
- Collection of existing metrics for IPv6 traffic; new link ME variables available for tracking IPv4 vs. IPv6 traffic
- Reporting (IPv6 Traffic Summary) to show how much IPv6 traffic is on your network in comparison to other types of traffic—IPv4 and non-IP.
- Graphing of IPv6 classes in Legacy and Sky UIs
- Host analysis of IPv6 hosts
- Display of IPv6 hosts in CLI command output: **traffic flow**, **traffic history**, **hostdb show**, **hostdb info**
- Classification testing of IPv6 flows
- Tracking of top talkers/listeners for IPv6 hosts

Feature Limitations

The IPv6 classification feature has the following limitations:

- IPv6 DNS lookup in hostdb is not supported.
- PacketShaper does not support compression or acceleration of IPv6 traffic.
- Synthetic transactions for IPv6 classes are not supported.
- You can do a packet capture of an IPv6 class, but you cannot filter for IPv6 hosts, IPv6 network ranges, host lists with IPv6 addresses, or IPv6 subnets.
- Only the last 32 bits of IPv6 addresses are sent in a flow detail record; because of this, multiple hosts could appear as the same IP address in IntellenceCenter Top Host reports.
- Host lists cannot contain IPv6 addresses.
- Classification by IPv6 host DNS names is not supported.
- You cannot specify a range of IPv6 addresses in a class matching rule.
- The host accounting feature does not track data for IPv6 addresses.
- The host adaptive response agents, such as High Bandwidth Hosts, do not monitor IPv6 hosts. Therefore, you should not use the host agents to track IPv6 hosts.
- While the PacketShaper does send IPv6 URLs to WebPulse, the current coverage of IPv6 websites is limited. As IPv6 websites grow in popularity, this coverage is expected to improve.

Licensing Overview

A Symantec license is required in order to use PacketShaper for its intended purpose: classifying and controlling network traffic. Note that a license is *not* pre-installed on PacketShaper S-Series appliances. Without a license, the appliance will act like a wire and pass traffic through it, without classifying, collecting data, or shaping the traffic.

There are two ways to install a license:

- **Option 1: Automatic License Retrieval**. If your PacketShaper has direct access to the Internet, the license key will automatically be synched to your appliance within 24 hours of running through the Guided Setup.

  > **Note:** Since the license synchronization may not be immediate, you can initiate the license retrieval process by issuing the " license update" on page 1235 CLI command.

- **Option 2: Manual License Installation.** When the PacketShaper does not have Internet access, you will need to manually install the license key. See " Install Initial License " on page 281.

To verify that your license is installed, check the *Licensing Status* on the Info tab in the the Sky UI or the Advanced UI. The *Licensing Status* indicates the date the license was activated and when it will expire. **show screen...**

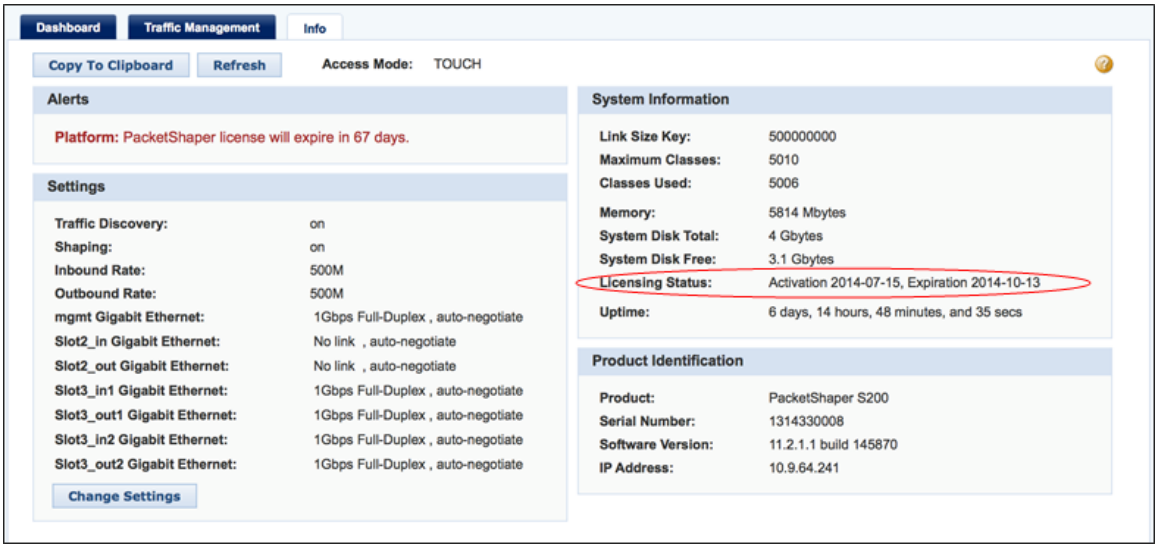

If a license will be expiring within the next 90 days, a message (such as *PacketShaper license will expire in 67 days*) appears in the CLI banner and on the Info tab in the Advanced and Sky UIs.

After purchasing or renewing a license, you can either wait for the Symantec licensing server to contact your PacketShaper (which it does periodically, every few hours or days), or issue the " license update" on page 1235 CLI command to initiate the process.

Matching Rules

What is a Matching Rule?

Matching rules define the criteria used by PacketShaper to identify traffic types. Every traffic class must have at least one matching rule. When traffic discovery creates a traffic class, it creates one or more matching rules to characterize the traffic type. In a similar fashion, when you create a class manually, you must specify the matching rules that describe the application's flows.

A traffic class can have multiple matching rules, which are treated as separate, distinct specifications. When PacketShaper tries to map a traffic flow to a class, it compares the flow with the criteria in the class' first matching rule. If PacketShaper does not find a match, it continues through the rules until a match is found or until it runs out of matching rules, in which case it moves on to the next class in the tree. If a specific traffic class cannot be found for a flow, the traffic is classified in the Default traffic class for the subtree.

The fields described here are found in the New Traffic Class and New Matching Rule windows. **show screen...**

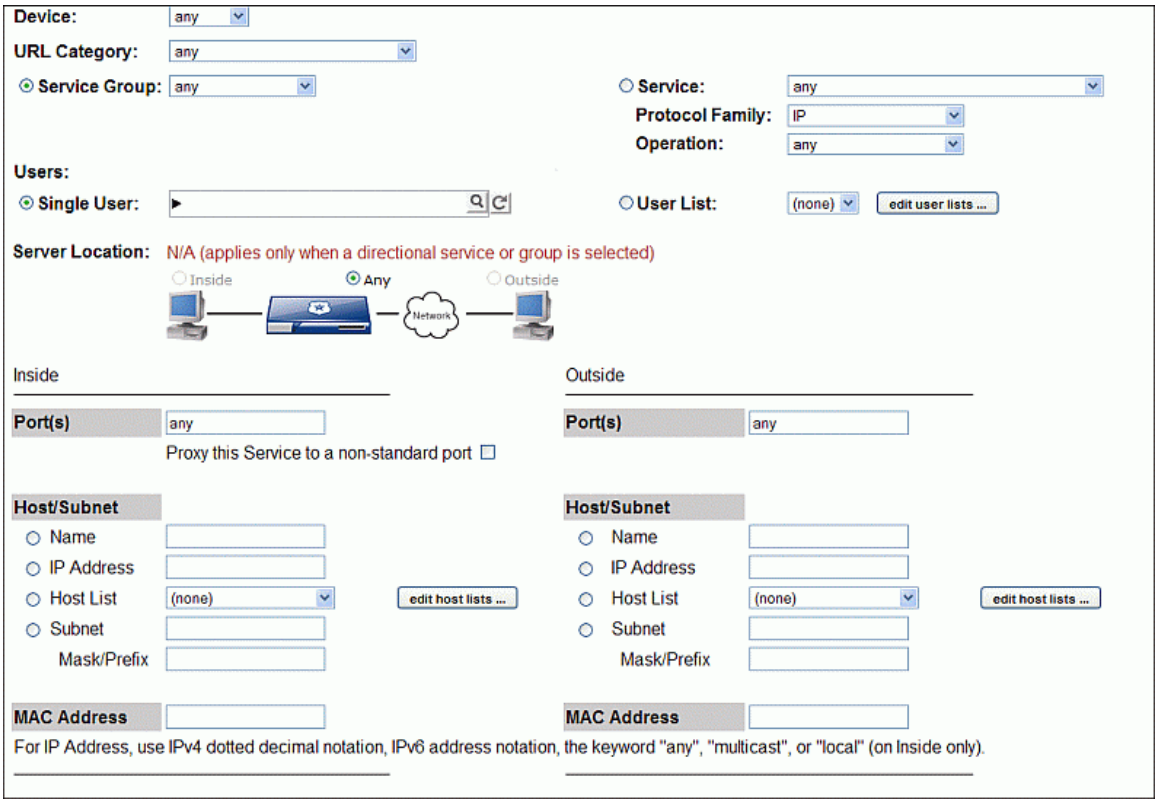

The following sections describe the criteria that can be specified in a matching rule:

Device

URL Category

Service Group

Service

Protocol Family

Operation

Users

Server Location

Host/Subnet

Application-Specific Criteria

Diffserv (DSCP) Matching Rule

VLAN Identification

Device

To match traffic that is passing through a specific interface card, you can create a class that specifies a physical port pair on the PacketShaper. For example, you can create a class that classifies all outbound traffic on Slot 2, into which a network interface card (NIC) is installed. You can then turn on class discovery for the Slot 2 class and PacketShaper will auto-discover all the applications, services, and protocols on this slot. The **Device** field lists only the installed interfaces on your unit.

> **Note:** If this unit is being managed by PolicyCenter, see "Device Matching Rules in PolicyCenter" on page 1039.

URL Category

To get visibility into the type of web traffic on your network, you can create or auto-discover URL category classes (Social_ Networking, Games, Gambling, and so forth). With this type of classification, the PacketShaper examines the URL, determines what category it belongs to, and then classifies the traffic into the appropriate category class.

- To classify by a particular URL category, choose the **URL Category** from the drop-down list. (Sky, Advanced UI)
- To create a class that discovers child classes for each URL category, choose **All Categories**, and then turn on class discovery (in Sky, Advanced UI) .
- To classify multiple categories in a single class, add a matching rule for each category. (Sky, Advanced UI)
- If you don't want to classify by URL category, keep the default setting (**any**).

Service Group

A *service group* is a set of services that have been grouped together based on a common functionality. For example, all services that classify peer-to-peer file sharing applications are in the P2P group. Service groups offer a quick way to classify and control a group of related services.

For a complete list of service groups supported by PacketShaper, see Applications, Protocols, and Services.

Service

The **Service** field lists all the services PacketShaper can identify for the selected protocol family. A service can be specified in a matching rule for IP, IPv6, DECNET, and Miscellaneous families. The service specifies an application or protocol for a traffic class. TCP and UDP can be specified here, although strictly speaking, they are not service protocols. The service name can be one of the built-in services or one of WebPulse's web application services.

> **Note:** Classifying by a web application requires WebPulse to be enabled. See Classify Web Applications.

For a complete list of services supported by PacketShaper, see Applications, Protocols, and Services.

Operation

The **Operation** field lists the web operations (such as Download_Video or Post_Messages) that PacketShaper can classify. The operation can be used in conjunction with a service to classify traffic for a web operation within a web application.

> **Note:** Classifying by operation requires WebPulse to be enabled. See Classify Web Applications.

Protocol Family

A protocol family represents a single network layer protocol and a number of higher-layer protocols and services.

PacketShaper can detect non-IP protocols only if your router is configured to handle these protocols.

The Protocol Family choices include:

Protocol Family	Classification Match
IP	Services on the IP network layer and TCP and UDP transport layers, for example FTP, HTTP (web), and Telnet.
IPv6	Services that can classify IPv6 traffic, for example CIFS-IPv6, FTP-IPv6, HTTP-IPv6.
IPX, SNA, AppleTalk, NetBEUI, DECnet, FNA	Non-IP protocol families.
others (miscellaneous)	Services for lesser-known non-IP protocols, such as Spanning Tree and Cisco Discovery. This protocol setting is treated like non-IP protocols, where host, port, service, and location are not applicable.
Any	Matches traffic for any protocol family.

Users

When PacketShaper is configured as a BCAAA client, you can manually create a class for a specific Active Directory user, group, or for a user list that has been defined in PacketShaper.

- Create a User Group Based Class (Sky UI, Advanced UI)
- Create a User-Based Class (Sky UI, Advanced UI)
- Create a Class Based on a User List (Sky UI, Advanced UI)

Server Location

For IP protocol traffic classes, the matching rule syntax uses the terms *inside* and *outside* to refer to an application's server location, relative to the unit. Typically, *inside hosts* are on your LAN and *outside hosts* are on the WAN or Internet.

> **Tip:** When setting Inside or Outside, answer the question: "If the chosen service uses a server, is the server found on the inside or outside of the unit?"

If you wish to capture traffic in a given class regardless of server location, select **any**. You will want to do this when creating a matching rule by client IP address or subnet. PacketShaper will create two matching rules one for inside and one for outside. Also, PacketShaper does not support the concept of a server for non-IP protocols NetBEUI, IPX, AppleTalk, SNA, DECnet, and FNA so these protocols do not have an inside or outside reference. For these protocol types, select **any** for the server location.

Classifying by Server Location

If you want to manage the traffic more precisely, create separate classes for the inside and outside server locations that is, /Inbound/HTTP/Inside, /Inbound/HTTP/Outside, /Outbound/HTTP/Inside, and /Outbound/HTTP/Outside. For example, you

may want to create these separate classes to be able to distinguish between uploads and downloads.

Example of Server Location

You can manually create traffic classes to account for an application's server location. For example, MySportsCompany, Inc. (a fictitious company) installed a unit on its network. The MySportsCompany web designer downloads images from the ESPN website.

The following table describes the traffic activity and the class that matches the activity:

Action and Traffic Detected	Traffic Class Match (direction/application/server_location)
1. A unit is deployed on the MySportsCompany LAN.	
2. MySportsCompany accesses the ESPN website an HTTP request. The ESPN web server is *outside*.	/Outbound/HTTP/Outside
3. ESPN web graphics are sent to MySportsCompany. The ESPN web server is still on the *outside* of MySportsCompany.	/Inbound/HTTP/Outside
4. The ESPN marketing staff wants to view the MySportsCompany website an HTTP request. The MySportsCompany web server is *inside*.	/Inbound/HTTP/Inside
5. The MySportsCompany web page is transmitted to ESPN. The MySportsCompany web server is *inside*.	/Outbound/HTTP/Inside

Ports

In the Inside and Outside **Port(s)** field, you can list the port or ports associated with the services for this class. In general, you should specify a port number *only* if you want to restrict classification to a specific port. Since many applications no longer limit transmissions to assigned port numbers, it is better to specify a service name.

You can define a service on a nonstandard port; however, the non-standard port must be a number that has not already been assigned. For example, in the matching rule, you can define the service as HTTP and the port as 8088.

For details about port number assignments, refer to the Internet Assigned Number Authority (IANA www.iana.org).

Port Number Range

To specify a range of port numbers for example, 5001 through 5005 use a dash (-) to separate the low and high values of the range (5001-5005).

Non-Contiguous Port Numbers

To specify non-contiguous port numbers for example, only ports 5001 and 5025 define separate matching rules within the same class.

Proxy this Service to a Non-Standard Port

A service proxy lets you identify applications running on ports that do not comply with the well-known port numbers defined by IANA. For example, you may know that a certain host (often a proxy server) is running a particular service on a specific non-standard port. In the matching rule, specify a combination of attributes, like host and port, and then use the **Proxy this Service to a non-standard port** option to associate the traffic with a known service.

Host/Subnet

Use the **Host/Subnet** matching rule field to isolate a single host or a range of hosts perhaps the network administrator's PC or a group of high-priority users.

You can specify source and destination hosts in a number of ways: by name, IP address, a range of IP addresses, a list of names, or subnet address and mask.

Name

If you have configured a DNS server in PacketShaper, you can specify a domain name in a matching rule. Also, if you configured a default domain name, you can use non-fully qualified domain names. If the name does not exist when you create the matching rule, PacketShaper warns you about the unknown name.

PacketShaper refreshes DNS name lookups in matching rules whenever they expire, according to the life that is returned in the lookup. This is configured separately at each name server. This client behavior is compatible with dynamic DNS.

IP Address

You can specify an IP address using one of the following ways:

- Enter a single IPv4 or IPv6 address for example, *207.78.98.254* or *2001:db8:1234:5678::2*.
- Specify a range of IPv4 addresses by separating the first and last addresses with a dash for example, *10.10.10.10-10.10.10.20*.

  > **Note:** When a parent class has an Inside IP address range and its child class has an Outside IP address range different from the parent range, you need to specify the parent's Outside IP range in the child class' matching rule. (The range is not inherited.) If you neglect to input the parent's range, a matching rule incompatibility error will occur.

- For multicast services, enter the keyword **multicast** to specify Class D IPv4 addresses (224.0.0.0 through 239.255.255.255) for the destination side that is, inside for inbound classes, outside for outbound classes.

  The classification process can handle broadcast and other multicast traffic if you specify a particular multicast address in

the matching rule. In this case, PacketShaper assumes that, since it is bridging this traffic to remote sites, it needs to count this traffic.

- Use **any** to indicate any host that is, no particular address.

> **Note:** When specifying an IPv6 address, make sure the service for the class is an IPv6 service. Likewise, when specifying an IPv4 address, the service must be IPv4.

Host List

A *host list* contains IPv4 addresses, DNS names, and/or subnets that traffic class matching rules can reference. Host lists let you specify many hosts in a single traffic class matching rule. The addresses in a host list do not need to be contiguous. See Create a Host List (in Sky, Advanced UI).

Subnet and Mask

Specify an IPv4 subnet mask to mask off network address bits, thereby defining a subset of hosts in a network for example, *255.255.255.0*

MAC Address

Enter a MAC address for IP or non-IP protocols to enable classification of traffic to and from a known host. The address must be entered in six-octet format for example, *08:12:34:56:A1:5C*. Note that you cannot classify Ethernet broadcasts (*FF:FF:FF:FF:FF:FF*).

When the host resides outside the unit, you must specify the MAC address in the "Outside" column. If the host is on the inside of the unit, use the "Inside" column. If you wish to match all traffic from a host, regardless of whether it is inside or outside of the unit, you need to create two separate matching rules.

In watch mode, you need to create two matching rules to classify by MAC. Even if you know the MAC resides on the outside, you need one matching rule with the MAC address on the Inside column and the other matching rule with the MAC address on the Outside column. You would first create the class based on an inside host, specifying the router's MAC address; you would then add a matching rule for an outside host and specify the same MAC address. When the PacketShaper is inline, you need only one matching rule.

The classification by MAC address feature allows you to create classes for traffic going through a specific router and/or router interface. For example, suppose your PacketShaper is connected to two WAN routers. With this feature, you can create a separate class for each router by specifying the router's interface MAC address in the matching rule. You can then assign a partition to each class to control bandwidth usage of the link.

Application-Specific Criteria

The **Criterion** field near the bottom of the *New Matching Rule* and *New Traffic Class* windows is available for certain types of traffic, specifically Citrix-ICA, DCOM, DICOM, FTP-Data-Clear, FTP-Data-Clear-IPv6, HTTP, HTTP-IPv6, HTTP-Tunnel, HTTP-Tunnel-IPv6, ICMP, ICMPv6, NNTP-Clear, Oracle-netv2, PostgreSQL, ProxySG-ADN, VProxySG-ADN-IPv6, RTCP-

I, RTP-I, SMTP-Clear, SMTP-Clear-IPv6, SOAP-HTTP, SSL, SSL-IPv6, and WAP. This field allows you to further differentiate these types of traffic. For example, web traffic can be differentiated by host DNS name or address, URL, content type, or web browser (user agent).

For details, see Application-Specific Matching Rule Criteria.

Diffserv (DSCP) Matching Rule

Diffserv (differentiated services) is a protocol for specifying and controlling network traffic so that certain types of traffic get precedence. *DSCP* (Differentiated Services Code Point) is a six-bit field in the IP header that specifies the per hop behavior for a given flow of packets.

The Advanced UI lets you create a class based on a Code Point or COS/TOS value (see descriptions below). The Sky UI lets you create a class based on DSCP.

Diffserv Type: Code Point (DSCP)

The Type of Service (TOS) field in the IP header is known as the Differentiated Services field, as defined in the Differentiated Services specification (RFC 2474). This field is divided into the following subfields:

- Differentiated Services Code Point (DSCP) (6 bits)
  Enter a value from 0 through 63.
- Currently unused (CU) (2 bits)

See Classify Diffserv Code Point Traffic.

Diffserv Type: COS/TOS

Many routing protocols make routing decisions based on the Type of Service (TOS) field in the IP header. This field is an indicator of the quality of service that is expected.

The TOS field contains:

- IP Precedence subfield, also known as the Class of Service (COS) sub-field, which is a 3-bit field
- Type of Service (TOS) subfield, which is a 4-bit field
- Unused bit that is set to zero

Applications set this field to tell routers how to prioritize packets. For example, weighted fair queuing (WFQ) algorithms in routers use this information.

See Classify Diffserv COS/TOS Traffic.

You can tell PacketShaper to match on these IP precedence bits during classification of IP protocols, then apply specific policies to manage these traffic types. For example, you could apply a policy that substitutes a different precedence value so that you can control the packet's priority when it reaches the router. For details on using a policy to substitute an IP precedence value, see Substitute Diffserv Values.

In the **IP Precedence (COS)** field, specify the IP precedence value to match. Enter a precedence value from 0 to 7 (where 7 is the highest precedence), or a range of values.

PacketShaper can check the **Type of Service (TOS)** field in the IP header to match on a service level for an application. Select one of the following values to match:

8 = minimize delay

4 = maximize throughput

2 = maximize reliability

1 = minimize monetary cost

0 = normal service

You might match on this TOS field, then use a policy to change the value.

VLAN Identification

PacketShaper has the ability to classify application traffic within multiple virtual LANs on a single Ethernet segment (an 802.1q trunk). For a multiple VLAN environment (where packets passing through the unit will have more than one VLAN header), only the outermost VLAN header is used for classification. The other headers are ignored.

You can create a VLAN identification matching rule that is based on:

- a single VLAN ID (a number between 0 and 4095)
- a range of VLAN IDs (for example, 100-200)
- any traffic flow that has a VLAN ID (specify **any** instead of a specific number)

To set up a VLAN identification matching rule, click **VLAN** when creating a class or matching rule in Sky or Advanced UI. Then enter the VLAN identification value.

Partition Overview

A *partition* manages bandwidth for a traffic class' aggregate flows, so that all of the flows for the class are controlled together as one. For information on managing individual flows, see Policy Overview.

For details on creating partitions, see:

Create Partitions in the Sky UI
or
Create a Partition in the Advanced UI

Managing Bandwidth

You can use partitions to:

- Protect mission-critical traffic by guaranteeing that a traffic class always gets a defined amount of bandwidth
- Limit aggressive, non-critical traffic by allowing that traffic class to consume only a defined amount of bandwidth
- Divide capacity
- Assign bandwidth dynamically to users
- Oversubscribe your link

Protecting Traffic

Partitions protect traffic by guaranteeing a defined amount of bandwidth for your mission-critical traffic classes.

For example, you could set a 128 Kbps partition for SNA traffic. This partition ensures that SNA will always have at least 128 Kbps of bandwidth available. Unpredictable surges of competing traffic will not interfere with SNA traffic.

Limiting Traffic

Partitions limit less important traffic by putting a cap on the amount of bandwidth a traffic class can use.

For example, say you have a 128K link. You can assign a 64 Kbps partition to FTP traffic. This prevents FTP traffic from consuming your entire link and blocking more important traffic (like Oracle or Citrix).

Another example of limiting traffic is restricting how much bandwidth a Class C subnet can use, regardless of how many sessions are active. You could create a 256 Kbps partition, burstable to 512 Kbps. The subnet's traffic would always get at least 256 Kbps, and could use as much as 512 Kbps if that bandwidth is available.

Dividing Capacity

Some traffic, such as Voice over IP (VoIP), requires a certain amount of bandwidth in order for the transmission quality to be acceptable.

For example, you could create a traffic class for VoIP traffic. A partition for the VoIP traffic class manages the aggregate VoIP traffic and the concurrent flows for the class. You could then combine the partition with a rate policy that defines a minimum rate for each flow.

In this example, by combining a rate policy with a partition, you can ensure that VoIP always has enough bandwidth to support the multiple flows during a VoIP session. Without this reserved chunk of bandwidth, the online conversation would be choppy and unintelligible, and the quality of service would suffer.

> **Note:** Partitions, like traffic classes, are one-way, so it is necessary to create a partition for each direction of a traffic type.

Assigning Bandwidth Dynamically

With dynamic partitions, subpartitions are created on the fly as users become active in a traffic class. This capability allows service providers or enterprise customers to guarantee a user a minimum amount of bandwidth at all times. This strategy is useful when a small portion of users will be active in any given time period.

For example, an ISP might have 20,000 customers, with up to 2,000 users logged on at a given time. With a dynamic partition established, subpartitions are automatically created for users as they log on. The dynamic subpartition allocates a minimum and maximum amount of bandwidth to each user. To accommodate new users when the maximum number of users is reached, the oldest non-active subpartition is removed so that bandwidth can be released to active users.

Oversubscribing

While not generally recommended, PacketShaper allows you to oversubscribe a partition that is, create partitions that allocate more total bandwidth than your link can provide. PacketShaper adjusts partition sizes to accommodate changing demands. For examples of this concept, see Hierarchical Partition Examples.

Initial Configuration

When you complete Guided Setup, the initial configuration provides a partition for the built-in /Inbound traffic class and a partition for the built-in /Outbound traffic class. Both /Inbound and /Outbound partition sizes correspond to the link bandwidth. You may also create additional partitions to divide the link bandwidth into multiple pieces. The number of partitions you can define depends on your model. For details, see Configuration Limits.

Types of Partitions

There are two types of partitions you can create: static or dynamic.

A *static partition* manages bandwidth for all flows within a particular traffic class. Static partitions can be fixed or burstable.

- A *fixed partition* allows an aggregate traffic class to use a defined amount of bandwidth, if needed. A fixed partition not only ensures that a specific amount of bandwidth will be available, but it also limits traffic to that same level.
- A *burstable partition* allows an aggregate traffic class to use a defined amount of bandwidth, and also allows that traffic class to access additional unused bandwidth, if needed. You can put a cap on a burstable partition, allowing it to access up to a maximum amount of bandwidth, or you can allow a burstable partition potentially to consume all available bandwidth.

In situations where you want to apply bandwidth limits to individual users, you can establish dynamic subpartitions for the traffic class. A *dynamic partition* carves up a static partition's bandwidth, creating subpartitions on the fly for new users. Subpartitions are children of a static partition.

Accessing Available Bandwidth

If a partition doesn't need all of its reserved bandwidth, only the *excess* unused bandwidth is available to be given out as *guaranteed rate* to other traffic. The reserved bandwidth can only be given out to other traffic as *excess* rate. For example, suppose the TCP class has a partition size of 100k with a limit of 400k and there is currently no TCP traffic. Up to 400k can be used by other traffic, but only 300k can be given out as guaranteed rate. The reserved bandwidth (100k in this example) cannot be given away as guaranteed rate.

When a partition needs its reserved bandwidth but other traffic is using it, the necessary bandwidth is returned to the partition and bandwidth usage by other traffic shrinks proportionately.

Hierarchical Partitions

All partitions are defined as *hierarchical* that is, partitions can contain partitions. For example, you might define a fixed partition for FTP and then create burstable child partitions for user groups that use this application.

This hierarchical approach enables application management for multiple groups, while controlling the group as a whole. For example, an ISP can subdivide a subscriber's partition with child partitions for each of the subscriber's departments. Likewise, an enterprise can allocate different amounts of bandwidth for a particular application to different groups of users. For instance, an enterprise might have a PeopleSoft class with a fixed partition. Within this partition, Human Resources could have a large burstable child partition and Accounting could have a smaller burstable partition.

> **Note:** The key concept for hierarchical partitions is that child partition minimums are limited to the parent partition minimum. When the sum of the minimum size of the child partitions exceeds the minimum size of the parent partition, the parent is oversubscribed and the child partitions are scaled proportionately. See the following examples.

Hierarchical Partition Examples

An ISP could set up one partition per customer (by server name or IP address) to match the customer's Service Level Agreement (SLA). The partitions enable the ISP to enforce SLAs for customers who consistently burst beyond their limits. In the following example, the FTP and HTTP classes (both 384 Kbps minimum) have oversubscribed their parent partition (512 Kbps minimum) by 256 Kbps. When a partition is oversubscribed by its child partitions, all partitions in its subtree are scaled accordingly. In this example, both child classes will be allocated 256 Kbps so that the sum does not exceed the parent's minimum size.

```
Inbound
 Customer1 (minimum partition size = 1 Mbps)
 Customer2 (minimum partition size = 512 Kbps)
 FTP (minimum partition size = 384 Kbps)
 HTTP (minimum partition size = 384 Kbps)
 Default
```

In the next example, the sum of the minimum size for the FTP and HTTP partitions is less than their parent, Customer2. In this case, both child classes will get at least 256 Kbps and can burst up to the size of Customer2 (1024 Kbps). Remember that if a partition is burstable and bandwidth is available, the partition can access the available bandwidth if needed.

```
Inbound
 Customer1 (minimum partition size = 512 Kbps)
 Customer2 (minimum partition size = 1024 Kbps)
 FTP (minimum partition size = 256 Kbps, burstable to 512 Kbps)
 HTTP (minimum partition size = 256 Kbps, burstable to 512 Kbps)
 Default
```

Policy Overview

A *policy* determines how an application's individual flows are treated in the context of competing applications, and allows you to manage bandwidth management on a flow-by-flow basis. With policies, you can give each flow of mission-critical traffic the bandwidth it needs for optimum performance, as well as protect it from greedy, less important traffic. In addition, policies can keep non-urgent traffic flows (such as FTP) from consuming more than an appropriate share of bandwidth.

Policy Types

PacketShaper offers the following policy types:

Priority	Establishes a priority for traffic without specifying a particular rate. Use priority policies for non-IP traffic types, or traffic that does not burst (for example, Telnet). For details: Set a Priority Policy in the Sky UI or Set a Priority Policy in the Advanced UI
Rate	Smooths bursty traffic, such as HTTP, using PacketShaper TCP rate-control technology. TCP rate control is an advanced congestion-avoidance mechanism whose goal is to prevent traffic from being sent at rates higher than the WAN connection, thereby greatly reducing queuing in router buffers and improving overall efficiency. For details: Set a Rate Policy in the Sky UI or Set a Rate Policy in the Advanced UI
Discard	Tosses all packets for a traffic class, thereby blocking the service. You might use this policy type for an application that is nonessential to your business and consumes too much of your network bandwidth. For details: Set a Discard Policy in the Sky UI or Set a Discard Policy in the Advanced UI
Ignore	Exempts a traffic class from bandwidth allocation and treats the traffic type as "pass-through" traffic. That is, the traffic won't be counted as part of the link traffic under management. Care should be taken when using this policy. If an ignore policy is placed on a class that is a major bandwidth consumer, other bandwidth allocation may be impacted. For details: Set an Ignore Policy in the Sky UI or Set an Ignore Policy in the Advanced UI

Restricts non-TCP traffic and intelligently rejects web and TCP traffic. Use this policy to redirect certain web users to alternate URLs. For details:

Set a Never-Admit Policy in the Sky UI
or
Set a Never-Admit Policy in the Advanced UI

Policies can be applied only to "leaf" classes that is, classes that do not have children. For example, a parent class, /Inbound/HTTP, may have child classes that differentiate one website from another: /Inbound/HTTP/ESPN and /Inbound/HTTP/MyCompany. In this example, policies can be placed on the ESPN and MyCompany classes (if they do not have children), but not on the parent class /Inbound/HTTP.

Never-Admit

If a class has a policy and then you create a child class, a Default class is automatically created in this subtree and the parent class policy is transferred to this class. If all child classes are later deleted, the policy is transferred back to the parent class.

See also:

Bandwidth Allocation

Policy/Partition Guidelines

Protect Critical Application Performance

Insulate Users of the Same Application

PolicyCenter Overview

Suppose a network manager installs a single PacketShaper on his company's network. He may spend one percent of his time updating the configuration of that single PacketShaper. This is not a large percentage of his work week, and so the addition of another four PacketShapers on the network (requiring an additional four percent of his time to configure and update) is not much more difficult for him to manage.

Now suppose that same network manager installs 95 more PacketShapers on the network. The effort that previously took just five percent of his time will now demand one hundred percent of his workday, leaving him time for little else except making every required change to a PacketShaper configuration 100 different times on 100 individual units.

What is needed is an economy of scale: a way to multiply the number of PacketShapers on a network without multiplying the amount of effort required to configure and maintain them. PolicyCenter is the solution, enabling network managers to manage many PacketShapers with the same amount of effort and time it takes to manage just a few.

> **Tip:** Refer to PacketGuide for PolicyCenter.

What is the Benefit of PolicyCenter?

PolicyCenter is a central management system that maintains multiple PacketShaper configurations on a single appliance. Because the configurations of all the PacketShaper units on the network are stored in a single place, they can be managed very efficiently. PolicyCenter can manage legacy PacketShapers (for example, PS 7500 and PS 12000) as well as PacketShaper S-Series appliances (PS-S200, S400, S500).

> **Tip:** Typically, you will manage the legacy and S-Series appliances in separate configurations since their settings vary and their image files are different.

Groups of PacketShapers can be assigned to a single PolicyCenter sharable configuration, allowing those units to operate with nearly identical configurations. When you change a configuration through PolicyCenter, the change immediately affects all units assigned to that configuration. It is this capability of PolicyCenter that truly provides the economy of scale: one single change to a PolicyCenter configuration can result in an instant configuration update on 100 PacketShapers.

PolicyCenter also allows you to:

- Deploy policies and partitions across multiple PacketShaper units .
- Distribute PacketShaper software upgrades and adaptive response action files. For legacy PacketShapers, PolicyCenter can also distribute customer portal and plug-in files.
- View a status summary of all managed PacketShapers.
- Monitor and manage the status of your unit and network with the adaptive response feature.

PolicyCenter Units Must be in Shared Mode

Individual PacketShaper units can be configured in either *local mode* or *shared mode*.

A unit running in *local mode* functions independently from other units, and has its entire configuration stored directly on its hard drive. Units configured in *shared mode* are assigned to a PolicyCenter configuration, and apply settings from their PolicyCenter configuration.

When a PacketShaper is in shared mode, PolicyCenter continually and efficiently synchronizes the unit's configuration on the PolicyCenter S-Series appliance with the configuration files on that unit's hard drive; therefore, if you switch from shared mode back to local, (or the network connection to the PolicyCenter S-Series appliance is lost) the unit's configuration in local mode will be the same as its last configuration in shared mode. Units in shared mode may be returned to local mode at any time.

RTM Overview

PacketShaper's response time measurement (RTM) enables you to measure the quality of users' network experience. RTM provides performance statistics, threshold monitoring, high-level problem indicators, and performance graphs. You can use the response-time analysis to highlight problems and confirm the effectiveness of solutions.

RTM Functions

The RTM feature performs the following functions:

- Tracks delay statistics for connection-based TCP traffic classes (for example, you can measure response times for applications, web traffic, individual hosts, or subnets). See Monitor Response Times.
- Breaks each response time measurement into network delay, server delay, and total delay. See Show Transaction Delays.
- Identifies the clients and servers with the slowest performance records. See Enable Worst Client and Server Analysis.
- Allows you to set acceptability standards and track whether performance meets those standards. See Set Thresholds.
- Offers current and historical performance data in tables and graphs, as well as in a MIB (management information base). For details on downloading the MIBs, see "Use PacketShaper MIBs" on page 1150.

The PacketShaper is positioned in the corporate network to monitor and control all the traffic that passes through it. This gives PacketShaper an opportunity to provide accurate response time measurements at low cost. Since PacketShaper already handles and classifies every packet, it can calculate the time traffic spends traveling between a client and a server. Instead of polling for response data, PacketShaper notes response times as traffic passes through the unit.

Benefits

Response time measurement quantifies what has traditionally been subjective information. Performance measurement allows you to:

- Recognize performance problems before they impact business
- Assess the results of configuration changes, including changes to PacketShaper's policy and partition configuration
- Provide and monitor meaningful Service Level Agreements (SLAs)
- Feed strategies for bandwidth management, server balancing, and topology planning

How Response Times are Tracked

Traffic class response times are tracked if the associated traffic is TCP and conforms to recognized transaction behavior, including request/response pairs, push bits, MSS (maximum segment size), and so on. PacketShaper stores at least one day's worth of RTM statistics at one-minute intervals and at least one month's worth of data at one-hour intervals.

Response times are recorded only in the client-to-server direction. If there is no traffic in this direction, the RTM values recorded will be zero. Note that the Monitor Response Time screen in the browser interface and the **rtm show** command in the CLI list only classes that have RTM data.

Application response times involve bi-directional traffic. After all, it hardly makes sense to track a request without tracking its corresponding response. When looking at RTM statistics, you need to determine which statistics represent the activity you are trying to capture.

For example, suppose you want performance data for your branch office accessing web servers that are located either at corporate headquarters or other companies. This traffic is represented by the /Inbound/HTTP class and the /Outbound/HTTP class. In this case, the response times that tell the performance story are displayed in the /Outbound/HTTP traffic class because it represents client-to-server traffic. While the relevant RTM data is in the /Outbound/HTTP class, you would most likely apply a policy to the /Inbound/HTTP class, because it represents the download of data to the client.

Sky UI Overview

The Sky UI is an alternative user interface to the Advanced UI. It offers the following features:

Dashboard A graphical dashboard that allows insight into the applications that are consuming the most bandwidth on the network and current link utilization.

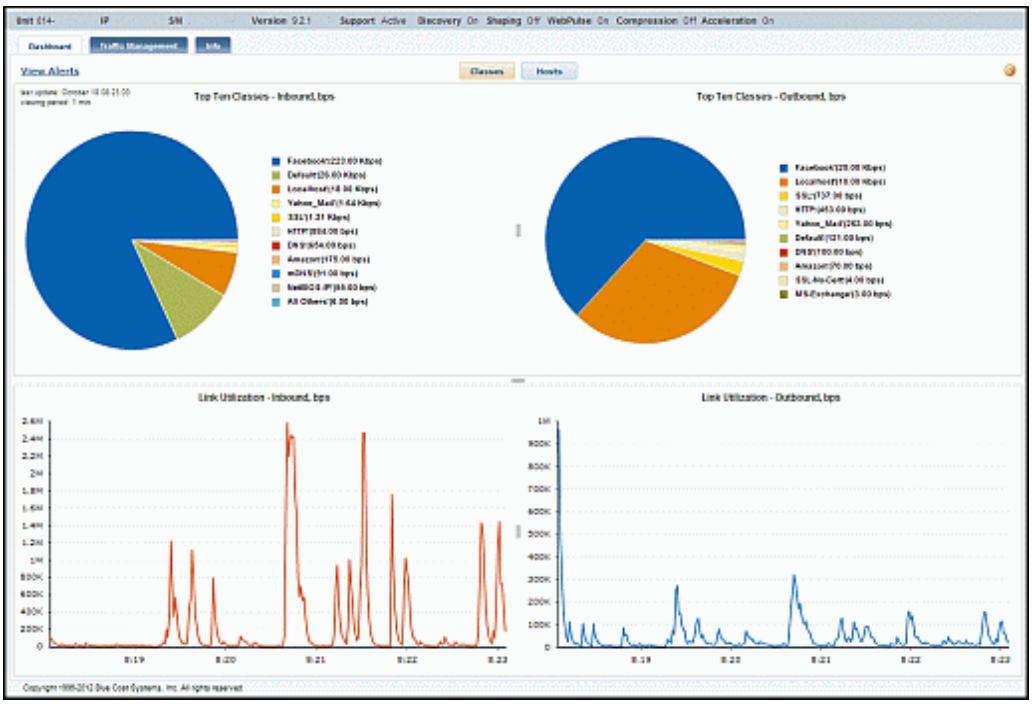

Class Tree Located in the top pane of the Traffic Management tab, the class tree displays a snapshot of the current activity for traffic classes, partitions, and policies. You can select one or more classes in the tree and then use the Reports module to create real-time or historical reports for those classes. Or select classes and use the Policy Manager to delete/modify the classes or apply policies to the selected classes.

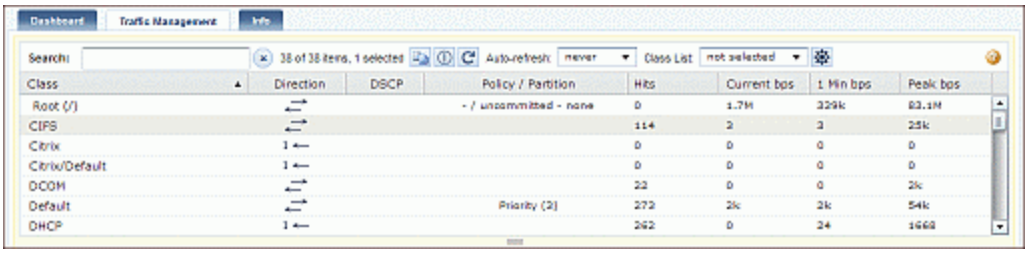

You can sort the class tree by any column, filter the view using the Search field, and create class list selections for future use. Note that the the Sky UI class tree displays differently than in the original UI: Inbound and Outbound classes are combined and the classes are sorted alphabetically.

Reporting A tool for creating real-time and historical reports

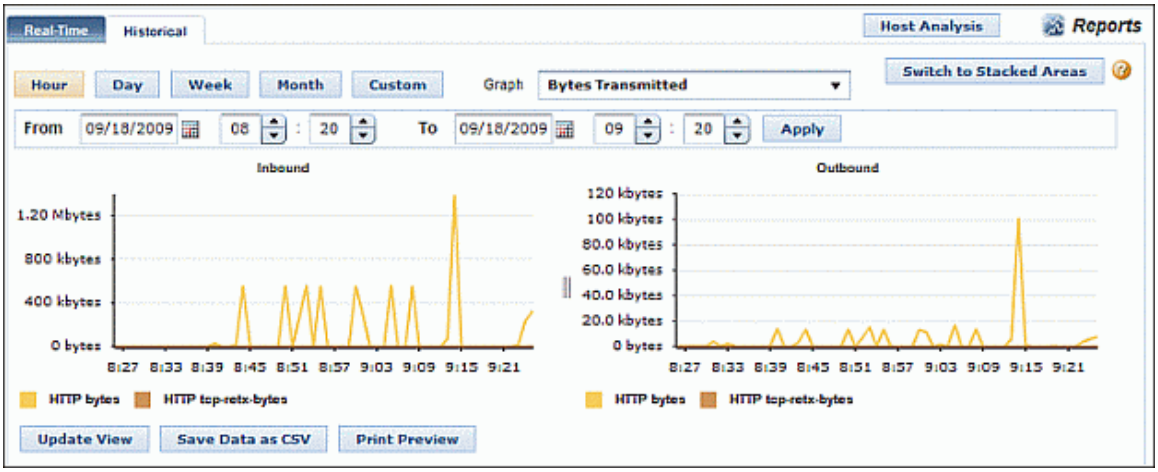

Policy Manager A simplified approach for creating traffic classes and controlling traffic with policies and partitions

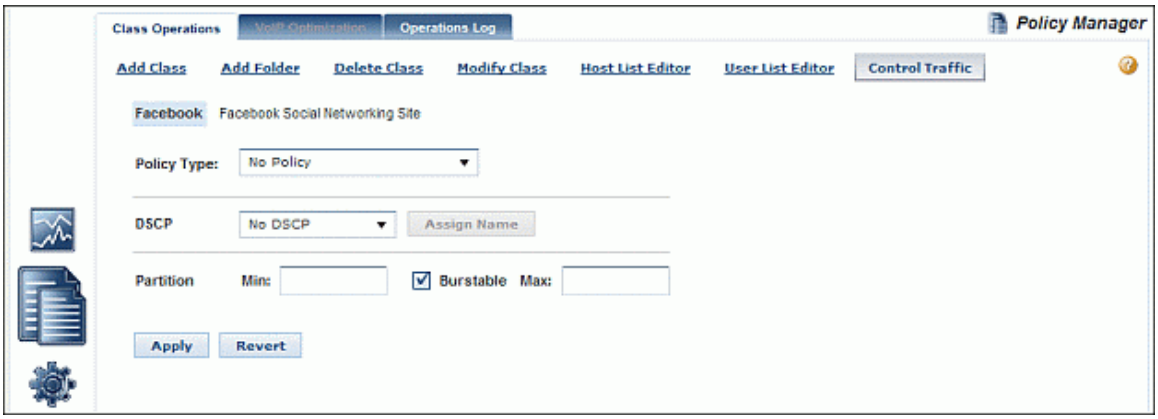

VoIP Optimization A VoIP Optimization tab that simplifies the configuration of VoIP traffic classes by recommending appropriate policy settings for each codec

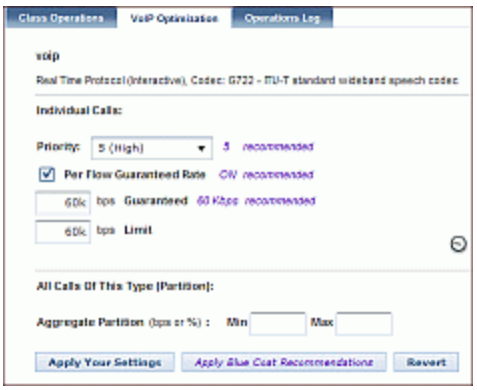

Advanced UI

Because not all PacketShaper features and capabilities are available in the Sky UI, the original full-featured user interface is still available: just click the **Advanced UI** link in the banner. the Sky UI in the current tab or window will be instantly replaced

with the original user interface.

Here is a partial list of tasks available only in the Advanced UI:

- Configure hardware and software features (Tip: The **Change Settings** button on the Info tab goes directly to the Advanced UI settings page)
- Create partition reports
- Schedule commands
- Perform top host analysis

To switch back to the Sky UI, click the **Blue Coat Sky** link in the banner. Note that you should not use the browser Back button for navigation between the two user interfaces, as it will display the login screen.

SNMP Overview

PacketShaper includes a Simple Network Management Protocol (SNMP) agent and PacketShaper enterprise Management Information Bases (MIBs). Using an SNMP management program, you can:

- Poll the PacketShaper for status information
- Read current values for the options described in the PacketShaper proprietary MIBs

You can configure PacketShaper to send unsolicited messages to any SNMP trap listener. PacketShaper trap messages report alarm conditions configured by default, for example a Power Supply Failure message is sent if the appliance's power supply fails. These are generally the same alarm conditions as might be reported at login to the browser interface on the Info page. See PacketShaper SNMP Traps for a list of supported traps.

PacketShaper supports SNMPv1, SNMPv2c, and SNMPv3.

- SNMPv1 and v2c rely on community strings for authentication.
- SNMPv3 provides greater security features for authentication, privacy, and access control.

Installation

SNMP agents are automatically installed on PacketShaper appliances. No additional re-configuration is required for the PacketShaper to function.

SNMPv1/v2c

By default, SNMPv1 and SNMPv2c are turned off until you set the look community string (password) and enable SNMP. You must set the look community string so that PacketShaper MIB information can be accessed by network management platforms and network management stations.

To set the PacketShaper read and write community strings via the browser interface, see Configure PacketShaper for SNMPv1 Support. See the command-line interface index to view the complete list of **setup snmp** commands.

SNMPv3

SNMPv3 provides a flexible and powerful framework for message security and access control. SNMPv3 authenticates packets with MD5 or the Secure Hash Algorithms (SHA), to ensure that a packet has not been tampered with in transit. This feature addresses the need for strong data origin authentication, unlike earlier versions of SNMP (v1 and v2c). SNMPv3's encryption feature provides additional security with Data Encryption Standard (DES) 56-bit encryption to hide the contents of an SNMP packet from unauthorized users. SNMPv3 also offers protection against replay attacks and message stream modification by checking timestamps and message stream integrity.

SNMPv3 encompasses a user-based security model for network management communications. Rather than using the two "read" and "write" community strings of prior SNMP implementations, administrators can create specific accounts for each

SNMP user and grant privileges through those user accounts. With SNMPv3, each user must belong to a group with a defined security model and security level, as well as a group access policy which specifies which SNMP objects the users can access, and which notifications the group's users will receive.

Model	Level	Authentication	Encryption	Model Features
SNMPv1 SNMPv2c	noAuthNoPriv	community string	No	Communication without encryption. Uses a community string match for authentication. This option provides no security, confidentiality, or privacy at all, though it might be useful for certain applications such as development and debugging.
SNMPv3	noAuthNoPriv	No	No	Communication without encryption or authentication. This option provides no security, confidentiality, or privacy at all, though it might be useful for certain applications such as development and debugging.
SNMPv3	authNoPriv	MD5 or SHA	No	Communication without encryption. Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Users must be authenticated before they access any of the values in the MIB objects on the agent.
SNMPv3	authPriv	MD5 or SHA	DES AES (128-bit key only)	Communication with both authentication and encryption. Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides AES or DES encryption, in addition to authentication. Users must be authenticated before they access any of the values in the MIB objects on the agent. In addition, all of the requests and responses from the management application to the SNMPv3 entity are encrypted, so that all the data is secure.

To use SNMPv3 with PacketShaper, you define values for users, access groups, views, remote users, and targets for notification. For details, see .Configure PacketShaper for SNMPv3 Support . You must also enable the SNMP functionality after defining the configuration.

Supported SNMP MIBs

Symantec is the authoritative source of the following MIB files, which can be downloaded from the Symantec download site.

BCSI-MIB.my	**Symantec Root MIB**
BLUECOAT-PACKETSHAPER-S-Series-MIB.my	**PacketShaper S-Series MIB** This MIB contains the variables that you can query on a PacketShaper. There are three tables for measurement data. The Link table contains measurement variables about each physical network interface on the PacketShaper. The Partition and Class tables contain measurement variables about each defined partition and traffic class, respectively. In addition to measurement variables, the MIB contains a number of predefined traps; see PacketShaper SNMP Traps for a description of the traps included in the Symantec MIB.

PacketShaper S-Series Response Time Measurement MIB for SNMPv1, v2c and v3	
BLUECOAT-PACKETSHAPER-S-Series-RTM-MIB.my	This MIB contains the measurement variables related to PacketShaper's Response Time Measurement (RTM) feature, including variables that track delay (total, network, and server), thresholds, and slow transactions. There are seven tables for response time configuration information or data about classes in the traffic tree. For more information about RTM, see " RTM Overview " on page 104.

PacketShaper also supports several variables in the following standard MIBs. The suggested source for these files is http://www.ietf.org.

RFC Number	RFC Title
RFC 1155	Structure and Identification of Management Information for TCP/IP-based Internets
RFC 1157	Simple Network Management Protocol Note that we do not return values for the Object Type egpNeighTable.
RFC 1212	Concise MIB Definitions
RFC 1213	Management Information Base for Network Management of TCP/IP-based internets: MIB-II We support following groups from MIB-2: ■ system ■ interfaces ■ address translation ■ ip (except for except the IP routing table) ■ icmp ■ tcp (tcpConnState is read-only) ■ udp (except for udpConnTable) Note that Symantec supported these groups in RFC 1156 before it was obsoleted by RFC 1213. RFC 1213 objects can be accessed with SNMPv1, v2, and v3 protocols
RFC 1215	Convention for defining traps for use with the SNMP
RFC 1493	Definitions of Managed Objects for Bridges. Symantec supports only dot1Bridge (dot1dBase, dot1dTp)
RFC1901	Introduction to Community-based SNMPv2
RFC 2011	SNMPv2 Management Information Base for the Internet Protocol using SMIv2
RFC 2012	SNMPv2 Management Information Base for the Transmission Control Protocol using SMIv2
RFC 2013	SNMPv2 Management Information Base for the User Datagram Protocol using SMIv2
RFC 2578	Structure of Management Information Version 2 (SMIv2)

RFC Number	RFC Title
RFC 2579	Textual Conventions for SMIv2
RFC 2580	Conformance Statements for SMIv2
RFC 3410	Introduction and Applicability Statements for Internet Standard Management Framework
RFC 3411	An Architecture for Describing SNMP Management Frameworks
RFC 3412	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
RFC 3413	SNMPv3 Applications
RFC 3414	User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)
RFC 3415	View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)
RFC 3416	Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)
RFC 3417	Transport Mappings for the Simple Network Management Protocol (SNMP)
RFC 3418	Management Information Base (MIB) for the Simple Network Management Protocol (SNMP). This document obsoletes definitions in RFC documents 1907.
RFC 3584	Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework

Note: PacketShaper does not support the Diffie-Hellman USM key exchange described in RFC 2786.

Standby Overview

In Symantec's high availability solution for PacketShaper S-Series, appliances can be installed in redundant network paths to provide link redundancy in case a component (PacketShaper, router, or switch) fails. This feature is called *Standby*, since the partner PacketShapers are on standby to take control if a link fails. In addition, the PacketShaper high availability solution can be applied in a network environment where asymmetric traffic tends to occur. All PacketShaper units in a configured high availability cluster are active and manage traffic on their own links, but they are simultaneously and continuously monitoring the health and status of its partners to make sure a unit has not gone down.

The example diagram below shows two PacketShaper S-Series appliances (different models) in a high availability group or *cluster*. Note that each PacketShaper manages its own link and has network routes to one another, but doesn't need to be directly connected by a cable to other partners.

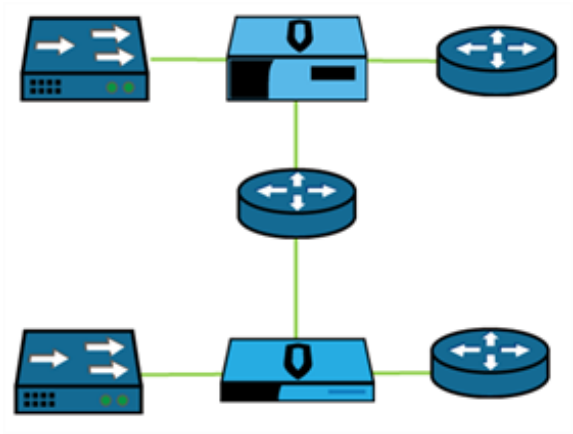

When a switch or router running redundancy protocol detects a link failure, it re-routes traffic to its alternate path. The PacketShaper in this alternate route will then take over traffic management responsibilities for the failed node, monitoring and shaping traffic that previously was managed by the PacketShaper in the disabled network path. This transition is seamless and transparent to users and administrators. The active node continues to manage its own traffic, in addition to managing the traffic on the disabled link.

In the diagram below, when *Switch A* detects that the other link is down, traffic is re-routed to *Switch A* (the switch designated in its redundancy configuration). In this example, *PacketShaper B* has failed, but if *Switch B* or *Router B* had failed, the same re-routing would occur. The traffic re-routing is indicated by the orange arrows in the diagram. *PacketShaper A* then takes over managing traffic that *PacketShaper B* was responsible for before the link went down. The green line illustrates that *PacketShaper A* continues processing its own traffic, in addition to that of the disabled link.

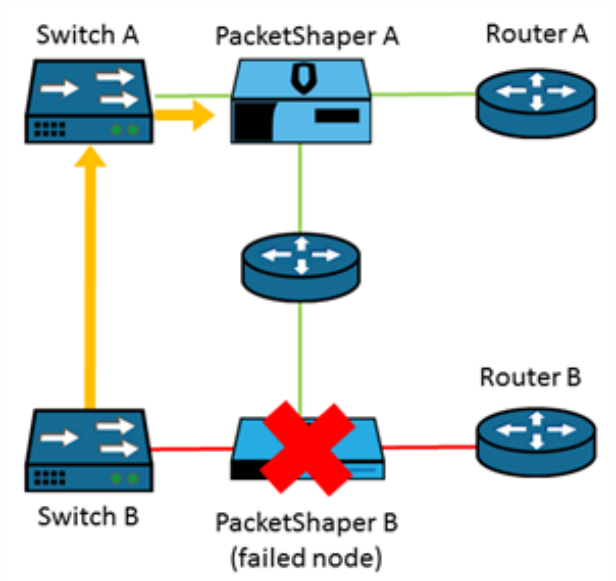

As soon as the other link is back up, traffic will no longer be re-routed to *Switch A*. *PacketShaper B* will resume traffic management of its own link, and *PacketShaper A* will return to managing only its own traffic.

Requirements

The standby feature is supported on all PacketShaper S-Series appliances running PS 11.6.1 or higher.

Caution: A high availability cluster of three or more partners requires additional PacketShaper memory than a cluster of two. The PS-S500 requires a minimum of 64GB of memory, the PS-S400 needs 32GB, and the PS-S200 requires 16GB. Some appliances may require a memory upgrade, available directly from Symantec.

- The PacketShaper standby functionality supports a five-partner cluster. Clusters with three or more partners require additional memory. (See note above.)
- The high availability cluster can comprise PacketShapers of any S-Series model—PS-S200, PS-S400, and/or PS-S500—with a variety of link key sizes. Each appliance must be licensed with a sufficient link size to manage its own traffic as well as the traffic of any of the partners in the cluster.
- The PacketShapers do not need to be directly connected (although they can be); they just need to be able to communicate with one another over the network with less than 50 ms network delay between each partner. They can be on different floors of the same building or even in different buildings, and there can be multiple hops between PacketShaper partners.
- Bypass must be set to open on all PacketShapers in the high availability cluster. When bypass is open, the PacketShaper becomes a point of failure in the network.
- Because bypass is open, the units should not be powered off when they sit in a standby configuration – doing so causes loss of connectivity on that link and all traffic is routed to the other path.

- For standby to work, each device must be able to communicate with the other device. If you have set the outside interface to list, you must add the partner's IP address to the Outside security list.

Additional Information

- The monitoring of partners in a high availability cluster has minimal impact on performance; the overhead for the feature is about 15-20 percent.
- When standby partners are down, the active and standby units share available bandwidth. Ten percent of available bandwidth is allocated to each standby device that is down. Thus, if one standby partner is down, the active system uses 90 percent of the available bandwidth and the failed device gets 10 percent. For each standby partner that is down, PacketShaper creates a partition with a size equal to 10 percent of the WAN link size.
- Asymmetric flows are supported.
- Standby partners communicate with each other using TCP/SSL connections over port 2014. Make sure this port number is open on your firewall and listed on any access lists.
- You can modify the configuration (for example, add classes or change policies) of a disabled PacketShaper, but when the PacketShaper resumes traffic management it does so with its original configuration; the temporary changes are discarded.
- When a PacketShaper is down in standby mode, the partner PacketShaper will merge the class tree of the failed PacketShaper. If the merged class name is greater than 127 bytes, the merged class name will be truncated.

Traffic Tree Overview (Advanced UI)

PacketShaper uses a hierarchical branching structure to organize and classify passing traffic, forming a *traffic tree*. The tree arranges *traffic classes* in parent, child, and sibling relationships, similar to folders and files in a standard file system. A traffic class identifies one type of traffic you want to manage together (an application, a protocol, a URL category, all traffic to one location, and many other possibilities). A class can have one or more subordinate child classes with more specific characteristics than its parent class.

For example, the structure of a traffic tree fragment might look like this:

- Inbound
 + SAP
 + FTP
 - Citrix
 + PeopleSoft
 + MsWord
 ...and so on...

This tree fragment indicates that traffic is first organized by travel direction (Inbound), then by application (SAP, FTP, or Citrix), and then sometimes by further refinement of on application (PeopleSoft over Citrix, MS Word over Citrix). SAP, FTP, and Citrix are sibling classes and they are all share the same parent class Inbound. When referencing a class, use a forward slash between each level, for example, */Inbound/Citrix/MsWord*. You can create up to 11 levels in a traffic tree.

With this same tree fragment, FTP traffic that is traversing a link from a WAN to a LAN would be classified in Inbound/FTP. Traffic from a Microsoft Word application running over Citrix MetaFrame would be classified in the Inbound/Citrix/MsWord class.

The traffic tree is displayed when you click the **Manage** or **Monitor** tabs.

The tree is automatically sorted from most to least specific; classes with more specific criteria are at the top of the tree. For example, a web class for a specific URL would come before the general HTTP class. Classes that are the same relative level of specificity are alphabetized. PacketShaper automatically determines the order of the classes in the tree and the only way you can change the order is to create an exception class.

Each traffic class has one or more *matching rules* that detail the criteria for determining which traffic matches the class. For each passing traffic flow, PacketShaper traverses the traffic tree, starting at the top, to look for a matching traffic class. It compares each class' matching rules with the characteristics of the traffic flow. If they don't match, the search continues, going to the next class on the list (the next sibling). If they do match and the class has no child classes, the search stops. If the class has child classes, the search continues with the class' children.

Any traffic not explicitly classified in a specific traffic class is classified as /Inbound/Default or /Outbound/Default.

Creating a Traffic Tree

When you run Guided Setup during initial configuration of your PacketShaper or when you reset the class tree, you can select from two different traffic trees:

- **Default**–This bare-bones traffic tree includes: /Inbound and /Outbound classes with a Default class for each, and a Localhost class for the inbound and outbound directions. You can build out this tree by turning on traffic discovery or by manually creating classes (Advanced, Sky) for the type of traffic you want to track.

- **Model**–A pre-configured tree that can be used as is, or modified to suit your needs. This tree organizes network traffic into folders of typical genres, such as VoIP, risky websites, business-critical applications and data, and recreational websites and applications. It includes classes based on service groups and URL categories.

The location of your PacketShaper (at a main site or at a branch location) and your management goals (tracking traffic on a per-application, per-location, or other basis) influence which type of traffic tree is best for you. For a description of the types of traffic trees and their respective pros and cons, see Traffic Tree Designs. For more on how to create a traffic tree, see Classify Network Traffic.

The maximum number of classes in your traffic tree is determined by your PacketShaper model. To look up the capacity of your model, see Configuration Limits.

Large Traffic Trees

If you have more than 2000 classes, there are several differences in traffic tree functionality:

- The **Manage** tab does not display the traffic tree you select the class you want to manage from a list. **show screen..**

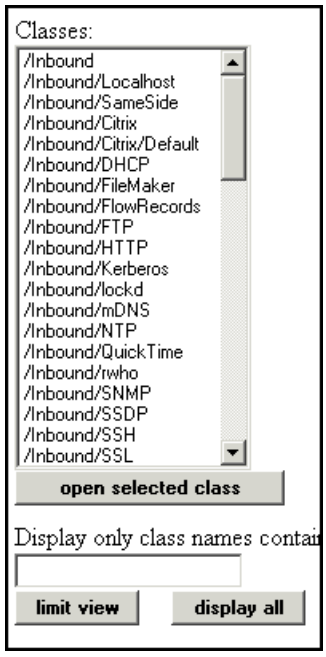

- The **Monitor** tab does not display the entire traffic tree – it shows 100 classes at a time, with **next** and **previous** links available for displaying other parts of the traffic tree. **show screen...**

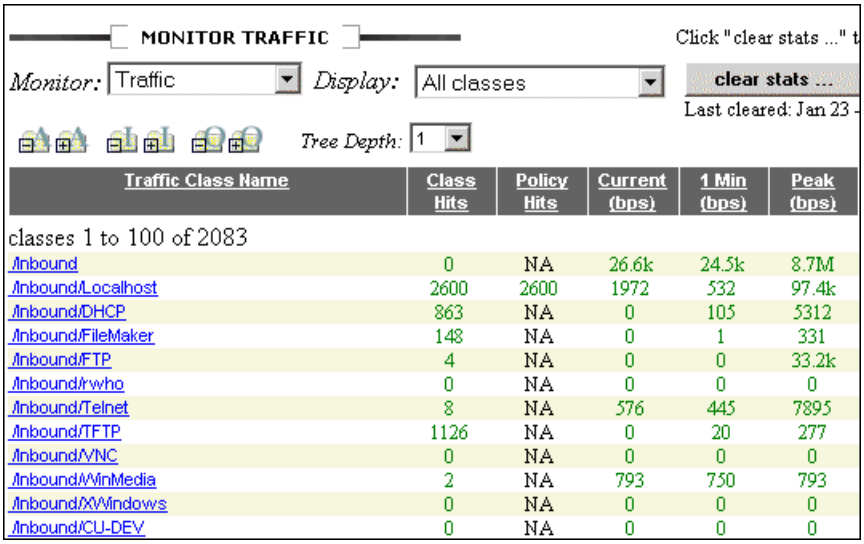

- Traffic class icons (described in the next section) are not shown on the tree.

These differences in functionality are necessary when working with large traffic trees. You will notice a delay as the browser initially loads the traffic tree, particularly on WAN links. However, once the tree is loaded into memory, the delay as you display each 100-class list will be minimal.

Traffic Class Icons

Next to each class name in the traffic tree is an icon, or a combination of icons. The icons represent the basic properties of a traffic class, such as whether a policy or partition has been applied and whether the class is a standard, exception, or default class. The following table describes each of the traffic class icons.

Icon	Description
	Standard Traffic Class A class that has one or more matching rules and is *not* a match-all class (see "Match-All Class" below)
!	**Exception Traffic Class** A class that has been tagged to sort above standard (non-exception) classes in the traffic tree. Exception classes give you the ability to redefine the search order that PacketShaper uses to find a match for a traffic flow. See Create an Exception Class.
	Policy A class that has a policy associated with it
	Partition A class that has a partition associated with it

Folder

A class with no matching rules. The folder is used to organize the tree structure and serves as a container only. It doesn't match any traffic, and cannot have a policy applied to it. To create a folder class, see Create a Folder Class.

Match-All Class

A class that typically appears at the bottom of a subtree. The bucket catches any traffic flows not matched by its sibling classes in the subtree.

A Default, match-all traffic class is created when the first child is added to a parent traffic class. For example, if you create a class for the Marketing subnet and then create a child class for a specific host on that subnet, a Default match-all class is created for the rest of the Marketing subnet. In this example, the default traffic class matches all traffic for the subnet (i.e., Protocol Family = any).

Icon Combinations

The traffic class icons can be combined together to designate a class with multiple characteristics. The following table provides examples of the icon combinations you might see on a traffic tree.

Standard Traffic Class with a Policy

A standard traffic class that has a policy assigned to it to control the quality of service for the traffic type.

Exception Traffic Class with a Policy

An exception traffic class that has a policy assigned to it.

PacketShaper automatically creates a Localhost class that is defined as an exception class with a priority policy set to a priority 6. This special class allows an administrator to always be able to access the unit in spite of congestion.

Default Match-All Class with a Policy

A match-all class, with the name Default, appears at the bottom of the /Inbound and /Outbound subtrees. These Default classes have inheritable policies. For details about inheritable policies, see Create an Inheritable Traffic Class.

> **Caution:** Symantec recommends that you do not change the default Priority 3 policy on the Inbound/Default and Outbound/Default classes.

Class Tree Overview (the Sky UI)

PacketShaper observes all the traffic running on the network and automatically compiles a list of the applications, protocols, and services being used, creating a *class tree* made up of *traffic classes*. A traffic class is a logical grouping of traffic flows that share the same characteristics – such as a specific application, protocol, IP address, or subnet. The process of creating traffic classes for observed network traffic is called *auto-discovery* and the process of categorizing the traffic into the appropriate traffic class is called *classification*. The class tree lists all of the classes that have been automatically discovered. You can manually create classes as well.

The Traffic Management tab provides an instant snapshot of the current activitiy for traffic classes, partitions, and policies. The class tree always displays in the top window pane; the content of the bottom pane varies. Use the icons on the left to select the area you want to use.

Icon	Name	Description
	Reports	Displays real-time and historical reports of the selected class(es)
	Policy Manager	Creates/deletes traffic classes, folders, policies, and partitions
	Settings	Controls display of the class tree, creates class lists, and displays log of CLI commands

Understanding the Class Tree

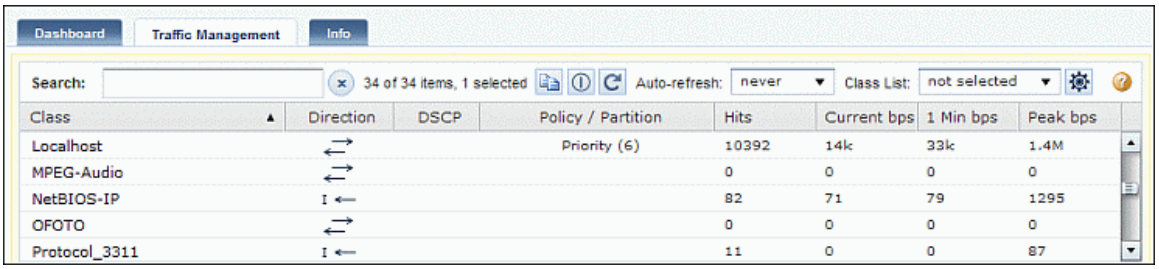

The class tree in the **Traffic Management** tab lists the following information for each traffic class:

Column	Description
Class	The name of the traffic class.
Direction	Indicates direction of traffic the class is monitoring: ■ Both directions ⇄ ■ Inbound only I ← ■ Outbound only O →

Note: The Direction column appears when the **Inbound & Outbound Combined** setting is enabled (as it is by default).

Column	Description
DSCP	The configured DiffServ Code Point (DSCP) policy for the traffic class. The DSCP can be a single value, a range of values, or a name (if one has been assigned to a DSCP value).
Policy	Shows the policy (if any) applied to the class. For a priority policy, the defined priority is shown. For a rate policy, the guaranteed rate setting, cap, and burstable priority level are shown.
Partition	The minimum and maximum amount of bandwidth allocated for a class subtree. *Uncommitted* is the difference between the current link speed and the sum of all committed partition minimums. *Uncommitted* indicates that the partition size is dynamically determined at run time and the minimum allocation is whatever is not committed to other partitions. *None* indicates that the partition is burstable without a limit. If the min and max values are the same, the partition is fixed and non-burstable. If a percentage appears after the partition size, the partition was specified as a percentage of the parent partition's size. Example: *450k (10%)* for the Inbound/HTTP class indicates that HTTP's partition size is 10% of the Inbound partition's size, resulting in a partition size of 450 Kbps (assuming the link size is 4.5 Mbps).
Hits	Number of times flows match the traffic class. Class hits occur only at the beginning of a flow or session. Classes created by traffic discovery do not show the first few connections for each traffic type. These initial hits are counted in the Default class. **Note:** In a configuration in which the PacketShaper unit sits between subnets and a router, traffic between two inside hosts will pass through the unit twice: once toward the router, then again back from the router to the other inside host. This results in double hit counts on a SameSide class. For example, four ping packets from one inside host to another would create eight class hits, as opposed to four with a standard setup (inside host to outside host).
Current bps	The current rate values in bits per second
1 Min bps	Last minute's average rate
Peak bps	The current peak values **Note:** Partitions have a 6% tolerance. Therefore, it is possible for an inbound traffic class to have a peak that exceeds the overall inbound peak.

Above the class tree are fields that allow you to perform tasks with the classes in the tree. See the following topics for more information:

URL Categories Overview

Web URLs can be grouped into various categories, such as social networking, gambling, pornography, news media, and shopping. PacketShaper is able to analyze URLs that users are requesting, determine what category the website belongs to, and classify the traffic into the appropriate category class. This gives you granular visibility into the type of web traffic on your network. You can also use PacketShaper's control capabilities to assign policies based on the category. For example, you can assign a never-admit policy to all the category classes with adult content. For a complete list of the URL categories PacketShaper can identify, see URL Categories.

The following steps and illustration describe PacketShaper's URL categorization and classification process. This process assumes the WebPulse feature is enabled and the traffic tree has been configured for category classes (see Category Classification below).

1. When a user requests a web page, PacketShaper first looks in its in-memory URL cache to see if it has already categorized the URL. If it finds the URL in the cache, it jumps to step 4. If not, it proceeds to step 2.
2. PacketShaper sends the URL to WebPulse, a cloud-based service provided by Symantec that (a) receives URL categorization requests and (b) responds with the ID* of the category associated with the URL.
3. This URL and its category ID are then stored in the URL cache in the PacketShaper system memory.
4. PacketShaper looks in the WebPulse map file for the category name associated with the URL's category ID.
5. PacketShaper classifies the traffic into the appropriate category class, and applies any policy that might be assigned to the class.

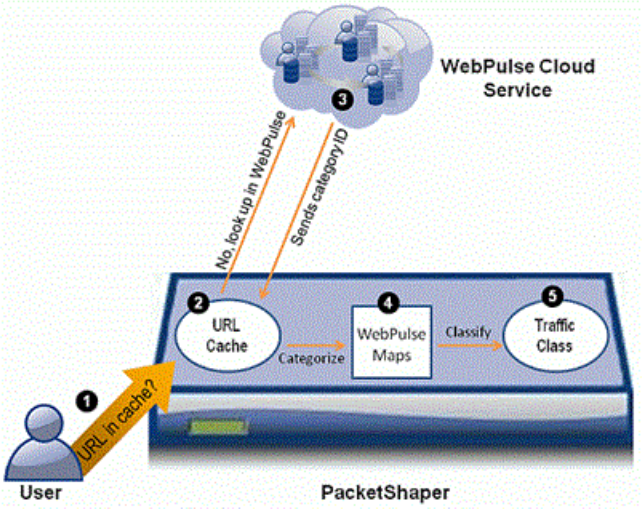

* This example assumes the URL has a single category ID. However, WebPulse can assign up to four categories to a given URL.

> **Note:** In addition to the URL category, WebPulse looks up the web application and operation. See WebPulse Overview for a description of the complete process.

WebPulse Express

With WebPulse Express, the URL database is stored locally (available on the PS-S500 appliance only). The following steps and illustration describe the process for identifying URL categories when the URL database is on box.

1. When a user requests a web page, PacketShaper first looks in its on-box URL database for the URL.
2. If it finds the URL in the database, it jumps to step 4. If not, it proceeds to step 3.
3. PacketShaper looks in its in-memory URL caches to see if it has already identified the URL. If it finds the URL in the cache, it proceeds to step 4. If not, PacketShaper sends the URL to the WebPulse cloud to identify the URL category.The URL and its category ID is stored in the URL cache in the PacketShaper system memory.
4. PacketShaper looks in the WebPulse map files for the category name associated with the URL's category ID.
5. PacketShaper classifies the traffic into the appropriate class, and applies any policy that might be assigned to the class.

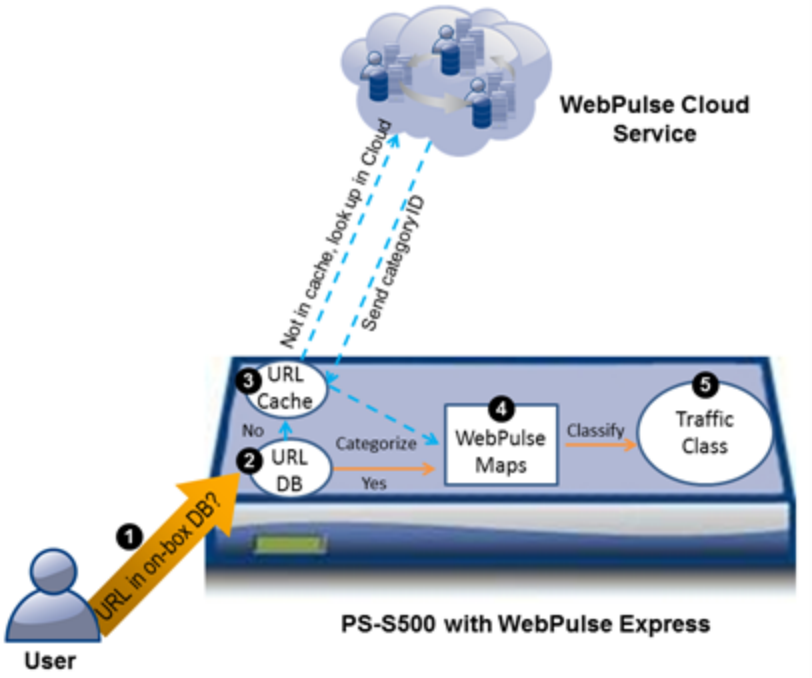

Category Classification

Each HTTP request is categorized; this means a flow may jump across multiple classes during its lifetime. When a user browses a website, such as cnn.com, the browser attempts to resolve all the embedded links on that site to make them viewable; this could result in many class and category hits for a single website visit.

SSL traffic is categorized based on the domain in the common name field of the SSL certificate.

When URL categorization is enabled, you will have the ability to:

- Create classes based on specific URL categories, and PacketShaper will then classify web traffic that corresponds to each of these categories into the appropriate class.
- Auto-discover category classes. Note that you cannot auto-discover category classes at the top level of the traffic tree.

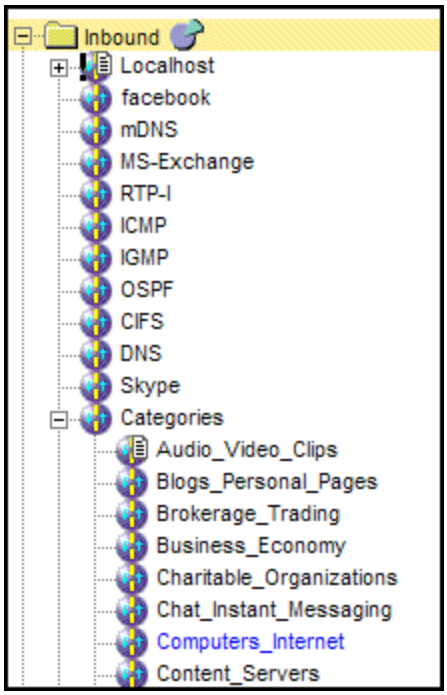

A traffic tree can have a mix of category-based classes and web-application classes. Web-application classes are more specific than category-based classes and will naturally sort to the top of the traffic tree, as shown on the tree to the left.

With a class tree that has both web-application and category-based classes, users can provide different treatment to a specific web application (such as Facebook) as opposed to other traffic in its URL category (Social Networking).

But what if you prefer Facebook to be classified into the Social_Networking category class instead of the Facebook web application class?

In situations where you don't need special treatment for a web application, you will want your category classes to get the class hits, instead of the web application classes. You can do this by making the Categories class an exception, as shown on the tree to the right. The exclamation point (!) indicates the class was made an exception.

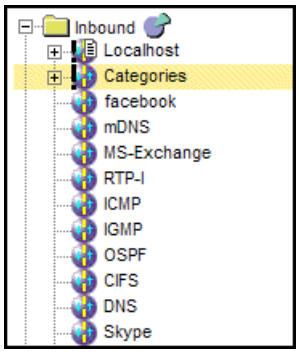

The WebPulse service can assign up to four categories to a given URL. For instance, a URL and its corresponding flow could contain the following category IDs:

catids[1 (Adult/Mature Content), 6(Nudity), 4(Sex Education)]

When such a flow is sent through the class tree, a class hit will occur if any one of the four categories matches a category defined in a class. This capability enables matching of nested categories within the class tree. For instance, if Farmville is categorized as both Social Networking and Games it will hit the Games class in the following class tree:

```
Inbound
 Social_Network
```

```
Games
Default
```

The above traffic tree would allow you to designate a different policy for social networking games than for other types of social networking traffic.

WebPulse Interaction

Symantec WebPulse contains a database of over 15 million entries, manages 85+ categories across 50 languages, and can return up to four categories per URL. URL categories are provided to the PacketShaper via requests to WebPulse service points located across the globe. Each service point is periodically pinged in order to ensure that categories are provided from the fastest service point.

When a user requests a URL (such as a brand new website) that has not already been categorized by WebPulse, the WebPulse dynamic categorization service queries the target website and retrieves key pieces of the page's content and context. After analyzing elements of the web pages, including the language used, the dynamic real time rating (DRTR) service determines both a likely category and a confidence factor that the rating is correct. Note that the PacketShaper will not hold up traffic during this process. In the rare event that the DRTR service does not yield a result with a high confidence level, the category rating request for the particular page is labeled unknown.

PacketShaper includes support for DRTR requests of plain HTTP traffic, but not HTTPS since only domain level information is available for SSL.

URL Cache

After WebPulse has looked up the category for a URL, PacketShaper stores the results in a URL cache in its system memory. Since serving from the cache is faster and more efficient than querying WebPulse, PacketShaper will always check the cache first before sending a query to WebPulse. WebPulse provides a “time to live” (TTL) for each categorization response. The PacketShaper scales this TTL by a multiple of seven to improve its cache efficiency. For instance, if the WebPulse service returns a TTL of one day for a particular URL, the PacketShaper will cache that URL and its associated categories for seven days.

The maximum cache size and number of entries in the cache varies by PacketShaper model. See Configuration Limits in PacketGuide.

The URL cache is stored in system memory, so the contents will be removed if the PacketShaper is reset or loses power. To preserve the cache contents, PacketShaper automatically backs up the cache to the 9.1026/urlcat/cache folder every 24 hours (replacing the previous backup). After a PacketShaper reboot, the most recent URL cache backup is automatically copied into system memory.

With WebPulse Express, the cache is used only if a URL is not in the on-box database and the URL needs to be looked up in the WebPulse cloud.

Feature Limitations

The URL categorization feature has the following limitations:

- Because the PacketShaper gives higher priority to flow delivery than to classification, it will never hold up flows to wait for a response from WebPulse. Therefore, the first few packets of a flow may get classified into a web or default class until WebPulse sends the URL category to the PacketShaper.
- Packet processing takes precedence over URL categorization. If the PacketShaper is under load, category requests may get queued, and some requests may be dropped.
- Behavior for asymmetrically applied redirect policies is non-deterministic for URL category-based classes since URL categorization is done out of path. Therefore, when applying never-admit policies with the redirect option, be sure to apply the policy to the category classes in both directions (Inbound and Outbound).

User Awareness Overview

PacketShaper has the ability to monitor and control traffic based on Active Directory user names and groups. You can create classes based on users and groups for the purpose of monitoring and controlling user-based traffic on the network. This functionality allows administrators to monitor top users of a service, monitor what services a specific user is using, or manage different policies for different users of the same service.

With the user awareness feature, you can:

- Display top users of a specific service (when Top Talkers or Top Listeners is turned on for the class). See Track Hosts and Users that Generate the Most Traffic.
- View the top ten users on the network (on the Top Hosts pie chart). See View the the Sky UI Dashboard.
- Report and apply policy on the user or group level (not just IP addresses)
- Create classes by user, group, or user list. See Classify Users (Sky UI, Advanced UI).
- Query and list hosts by user name on the Host Analysis page. See Perform Host Analysis.

Support for International Characters in User and Group Names

PacketShaper supports the display of international characters in user and group names. Whether your user and group names have simple accent marks over some characters or use foreign character sets such as Chinese or Arabic, the PacketShaper user interface is able to display these international user names in reports and matching rules.

Additional Information:

- To see UTF-8 (Universal Character Set Transformation Format 8-bit) characters in the CLI, you must configure your terminal client to display either UTF-8 or the specific character set you use.
- Although a user name matching rule can contain UTF-8 characters, the traffic class name cannot.

Requirements

The user awareness feature requires that you install the Symantec Authentication and Authorization Agent (BCAAA) version 6.0 on a Windows Domain Controller or other member server (running version Windows 2003 or 2008). BCAAA is a separate software program that converts IP addresses to user names so that PacketShaper can classify and report on network users and groups. When planning your BCAAA installation, you should install one instance of BCAAA per PacketShaper. Note that in certain situations, multiple PacketShapers can share a single BCAAA server; see Planning for your BCAAA Installation.

In addition, you need to configure BCAAA on the PacketShaper to establish communication between the two devices. You must enable BCAAA on the PacketShaper and specify the IP address and port number of the BCAAA server. See Configure PacketShaper as a BCAAA Client.

When BCAAA integration is disabled, the PacketShaper will use the latest user/IP mapping available stored in the PacketShaper cache.

> **Caution:** To prevent any impact on ProxySG authentication functionality, PacketShaper should not share a BCAAA server with a ProxySG.

Feature Limitations

The user awareness feature has the following limitations:

- See Configuration Limits in PacketGuide for the maximum number of Active Directory users and groups supported for each PacketShaper model.
- The **Traffic Class Test** command does not currently support the testing of user- or group-based classes.
- When a PacketShaper is managing traffic from multiple office locations within a company network, users may be seen as both source and destination IPs. If neither user (source or destination) is in the IP-user cache, the source user of a flow gets resolved and placed in the cache. However, if the source user is not cached but the destination user is cached, the source user in a flow will not be resolved.
- If Active Directory assigns an IP address to a different user while the original user-IP mapping is still in the PacketShaper cache, the stale mapping can cause incorrect user name identification in PacketShaper. The IP addresses that have a user mapping are stored in the cache for one hour by default; IPs with no user mapping are stored for 30 minutes. You can adjust the cache timeout settings on the System Variable setup page.

Watch Mode Overview

PacketShaper's *watch mode* is used when you don't want the PacketShaper to be cabled into the main data path when your situation requires that the unit monitor the traffic non-inline. This type of deployment is sometimes required for data centers that have restrictions on the introduction of inline elements into the network.

When a unit is in watch mode, it passively monitors the traffic on your network and performs all the traffic classification and reporting tasks that PacketShapers typically do: identify your network traffic, track bandwidth utilization, and report on the performance of your applications and network. It can provide you with insight into what applications are being used on your network and highlight performance issues. Consistent with its role as a passive watcher, a PacketShaper in watch mode will not perform traffic shaping; any existing policies and partitions will be ignored.

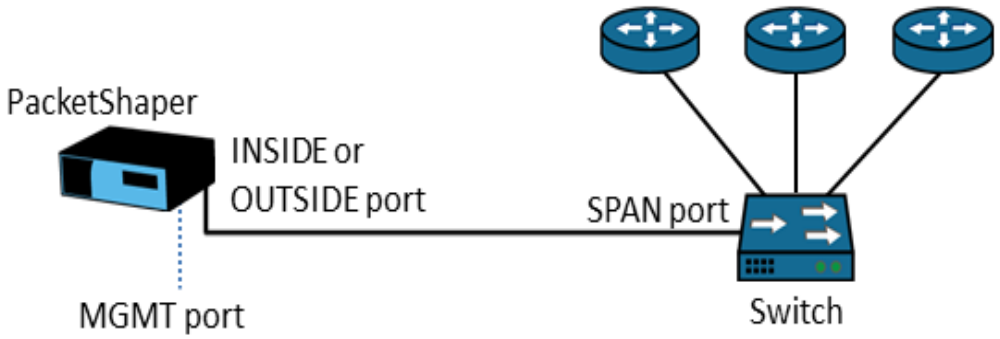

In watch mode, the PacketShaper's INSIDE or OUTSIDE port(s) are used for monitoring network traffic and the SYS MGMT port is used for managing the unit itself.

Watch mode is supported in a variety of network configurations. One example is shown above.

A PacketShaper in watch mode includes the following features:

- Is compatible with SPAN ports, mirrored switch ports, and hubs connecting multiple routers
- Can monitor traffic to/from up to 256 WAN routers
- Can receive traffic from multiple network segments (if NICs are installed)
- Can be managed from a separate network than the one being monitored
- Can be connected to the SPAN port via the PacketShaper's INSIDE or OUTSIDE port

Non-Inline Deployment Requirements and Limitations

PacketShaper's non-inline deployment has the following requirements and limitations:

- The host show command will not display output if both of the main NIC ports are disconnected and watch mode is enabled.

- Device classification cannot be used in watch mode. Because PacketShaper performs classification by device (NIC) when a packet exits the PacketShaper, and watch mode drops all packets, the packets do not exit on any interface and thus cannot be classified by interface. If you are using watch mode and want to classify and/or control traffic by device, you can create classes based on the MAC address of the watch mode routers.
- VoIP latency cannot be measured by a PacketShaper in watch mode.

Web Application Classification Overview

In addition to identifying the category of a URL, the WebPulse cloud service is able to inspect a URL and identify the type of web application and operation. PacketShaper can then use this information to classify the traffic into classes based on the web app or operation. By using the WebPulse cloud service, PacketShaper can expand its classification capabilities to include over 100 web applications, such as Facebook, Gmail, LinkedIn, Twitter, and Flickr and almost 20 different web operations (such as Download Attachment, Play Video, and Send Email). Thus, you can create the following types of classes:

- A specific web application for any operation (Service=Facebook, Operation=any)
- A specific operation for any web application (Operation=Upload Pictures, Service=any)
- A specific web application for a specific operation (Service=Facebook, Operation=Upload Pictures)

To see current lists of the supported web applications and operations, go to:

http://sitereview.bluecoat.com/applications.jsp
http://sitereview.bluecoat.com/operations.jsp

If traffic discovery is turned on and WebPulse is enabled, PacketShaper will automatically discover web applications, just as it discovers other services. The web application classes appear in the traffic tree, mixed in with other service-based classes. Operations are not auto-discovered; these classes must be manually created.

As you can with any type of class, you can use PacketShaper's control capabilities to assign policies based on the application or operation. For example, you can create a policy to block the playing of Netflix or YouTube streaming videos.

The following steps and illustration describe PacketShaper's process for identifying web applications and operations of URLs.

1. When a user requests a web page, PacketShaper first looks in its in-memory URL cache to see if it has already identified the URL. If it finds the URL in the cache, it jumps to step 4. If not, it proceeds to step 2.

2. PacketShaper sends the URL to WebPulse, a cloud-based service provided by Symantec that (a) receives web application/operation identification requests and (b) responds with the IDs of the application and operation associated with the URL.

3. This URL and its web application and operation IDs are then stored in the cache in the PacketShaper system memory.

4. PacketShaper looks in the WebPulse map file for the service name associated with the URL's application ID. It also looks up the operation ID to determine the action type.

5. PacketShaper classifies the traffic into the appropriate class, and applies any policy that might be assigned to the class.

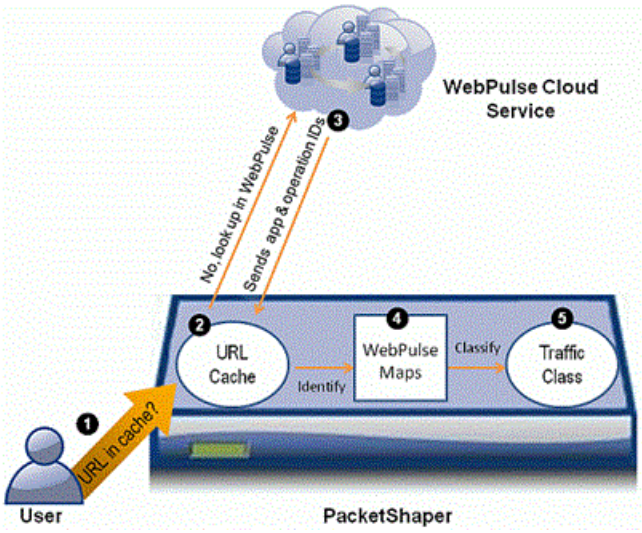

> **Note:** In addition to the web application and operation for a URL, WebPulse looks up the URL category. See WebPulse Overview for a description of the complete process.

WebPulse Express

With WebPulse Express, the URL database is stored locally (available on the PS-S500 appliance only). The following steps and illustration describe the process for identifying web applications and operations of URLs when the URL database is on box.

1. When a user requests a web page, PacketShaper first looks in its on-box URL database for the URL.
2. If it finds the URL in the database, it jumps to step 4. If not, it proceeds to step 3.
3. PacketShaper looks in its in-memory URL caches to see if it has already identified the URL. If it finds the URL in the cache, it proceeds to step 4. If not, PacketShaper sends the URL to the WebPulse cloud to identify the web application and operation associated with the URL.The URL and its IDs for web app and operation are stored in the URL cache in the PacketShaper system memory.
4. PacketShaper looks in the WebPulse map file for the service name associated with the URL's application ID. It also looks up the operation ID to determine the action type.
5. PacketShaper classifies the traffic into the appropriate class, and applies any policy that might be assigned to the class.

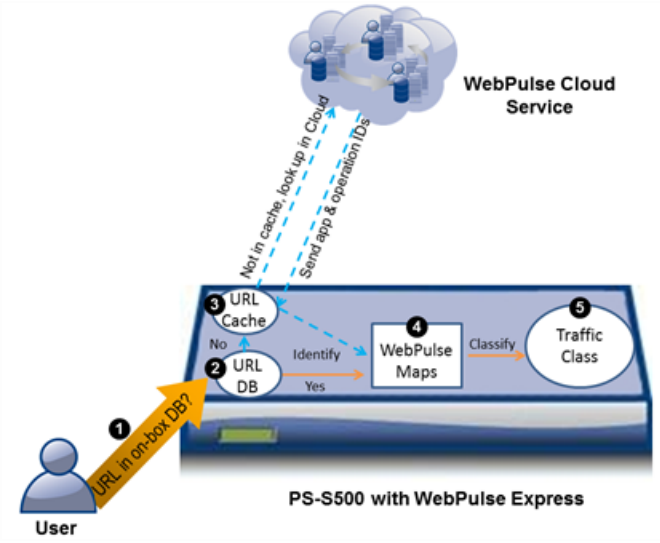

URL Cache

After WebPulse has looked up the web application and operation IDs for a URL, PacketShaper stores the results in the URL cache in its system memory. Since serving from the cache is faster and more efficient than querying WebPulse, PacketShaper will always check the cache first before sending a query to WebPulse. WebPulse provides a "time to live" (TTL) for each URL response. The PacketShaper scales this TTL by a multiple of seven to improve its cache efficiency. For instance, if the WebPulse service returns a TTL of one day for a particular URL, the PacketShaper will cache that URL and its associated web application and operation IDs for seven days.

The maximum cache size and number of entries in the cache varies by PacketShaper model. See Configuration Limits in PacketGuide.

The URL cache is stored in system memory, so the contents will be removed if the PacketShaper is reset or loses power. To preserve the cache contents, PacketShaper automatically backs up the cache to the 9.1026/urlcat/cache/ folder every 24 hours (replacing the previous backup). After a PacketShaper reboot, the most recent cache backup is automatically copied into system memory.

Legacy Built-In Web Applications

Previous versions of PacketShaper contained built-in services to classify a number of web applications, such as Facebook, YouTube, and WebEx. Starting in v9.2, these built-in web application services are replaced by those that WebPulse can identify when WebPulse is enabled. In addition, some of these web apps now have enhanced classification so that they can be further sub-classified into Operations (for example, Facebook/Upload_Pictures).

If you turn WebPulse off, PacketShaper returns to using the legacy built-in web application services for classification, as it had done in previous versions.

Feature Requirements

The web application classification feature has the following requirements:

- WebPulse must be enabled on the PacketShaper.
- PacketShaper must have access to the Internet.
- WebPulse Express requires a PS-S500 with more than 32 GB of RAM and a WebPulse Express subscription.

Feature Limitations

The web application/operation classification feature has the following limitations:

- Because the PacketShaper gives higher priority to flow delivery than to classification, it will never hold up flows to wait for a response from WebPulse. Therefore, the first few packets of a flow may get classified into a web or default class until WebPulse sends the web application/operation IDs to the PacketShaper.
- Packet processing takes precedence over web application identification. If the PacketShaper is under load, URL queries may get queued, and some requests may be dropped.
- Behavior for asymmetrically applied redirect policies is non-deterministic for web-application classes since this identification is done out of path. Therefore, when applying never-admit policies with the redirect option, be sure to apply the policy to the web application classes in both directions (Inbound and Outbound).
- Web operation classification is not supported for application traffic that is encrypted by SSL.

WebPulse Overview

Symantec products such as PacketShaper are able to send URLs to the WebPulse cloud service or an on-box URL database so that it can identify the URL category (social networking, gambling, pornography, news media, shopping, etc.), web application (such as Facebook, Gmail, LinkedIn, Twitter, and Flickr), and operation (such as Download Attachment, Play Video, and Send Email). PacketShaper can then use this information to classify the traffic into classes based on the URL category, web application, or operation. By using WebPulse, PacketShaper can expand its classification capabilities to include over 85 categories, 170 web applications, and 20 different web operations.

WebPulse contains a database of over 15 million entries and is growing every day. It provides URL identification to the PacketShaper via requests to WebPulse service points located across the globe. Each service point is periodically pinged in order to ensure that category, application, and operation IDs are provided from the fastest service point.

The following steps and illustration describe PacketShaper's process for identifying categories, web applications, and web operations of URLs.

1. When a user requests a web page, PacketShaper first looks in its in-memory URL caches to see if it has already identified the URL. If it finds the URL in the cache, it jumps to step 4. If not, it proceeds to step 2.
2. PacketShaper sends the URL to WebPulse to identify the category, web application, and operation associated with the URL.
3. The URL and its IDs for category, web app, and operation are stored in the URL cache in the PacketShaper system memory.
4. PacketShaper looks in the WebPulse map files for further information. For example, the category map file contains the category name associated with the URL's category ID.
5. PacketShaper classifies the traffic into the appropriate class, and applies any policy that might be assigned to the class.

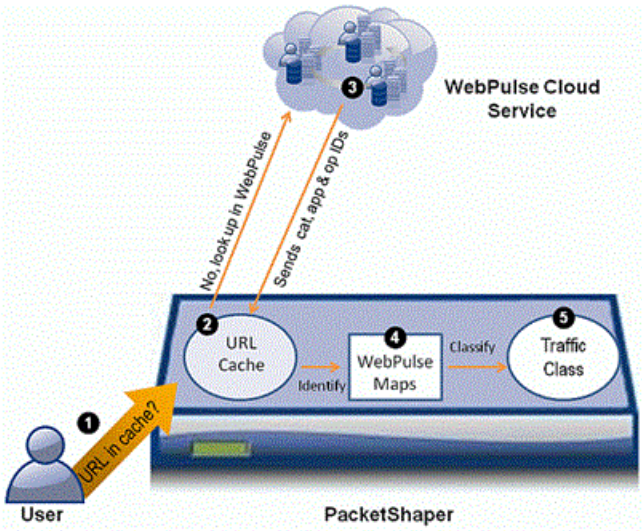

WebPulse Express

Caution: This feature is available only on PS-S500 appliances with 64GB of memory. See "setup webpulse express" on page 1367.

With WebPulse Express, the URL database is stored locally (available on the PS-S500 appliance only). The following steps and illustration describe the process for identifying categories, web applications, and operations of URLs when the URL database is on box.

1. When a user requests a web page, PacketShaper first looks in its on-box URL database for the URL.
2. If it finds the URL in the database, it jumps to step 4. If not, it proceeds to step 3.
3. PacketShaper looks in its in-memory URL caches to see if it has already identified the URL. If it finds the URL in the cache, it proceeds to step 4. If not, PacketShaper sends the URL to the WebPulse cloud to identify the category, web application, and operation associated with the URL.The URL and its IDs for category, web app, and operation are stored in the URL cache in the PacketShaper system memory.
4. PacketShaper looks in the WebPulse map files for further information. For example, the category map file contains the category name associated with the URL's category ID.
5. PacketShaper classifies the traffic into the appropriate class, and applies any policy that might be assigned to the class.

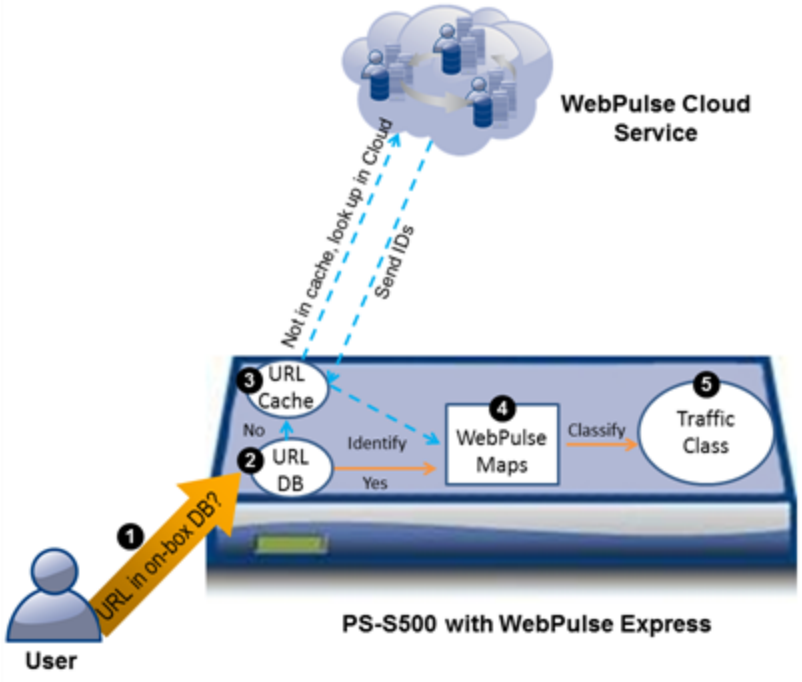

URL Cache

After WebPulse has looked up a URL in its database in the cloud, PacketShaper stores the results in a URL cache in its system memory. The URL and its category, web application, and operation IDs are stored in this cache. Since serving from the cache is faster and more efficient than querying WebPulse, PacketShaper will always check the cache first before sending a query to WebPulse. WebPulse provides a "time to live" (TTL) for each URL response. The PacketShaper scales this TTL by a multiple of

seven to improve its cache efficiency. For instance, if the WebPulse service returns a TTL of one day for a particular URL, the PacketShaper will cache that URL and its associated IDs for seven days.

The maximum cache size and number of entries in the cache varies by PacketShaper model. See Configuration Limits in PacketGuide.

The URL cache is stored in system memory, so the contents will be removed if the PacketShaper is reset or loses power. To preserve the cache contents, PacketShaper automatically backs up the cache to the 9.1026/urlcat/cache/ folder every 24 hours (replacing the previous backup). After a PacketShaper reboot, the most recent URL cache backup is automatically copied into system memory.

If WebPulse is disabled, PacketShaper does not look up URLs in the WebPulse cloud or in the URL cache.

With WebPulse Express, the cache is used only if a URL is not in the on-box database and the URL needs to be looked up in the WebPulse cloud.

Monitoring WebPulse-Related Traffic

If you want to monitor or control WebPulse query traffic, you can create a class based on the WebPulse service. There are two ways to do this:

- Manually create a WebPulse child class of the Localhost classes.

  or

- Turn on class discovery for the Localhost classes. (Localhost represents PacketShaper management traffic.)

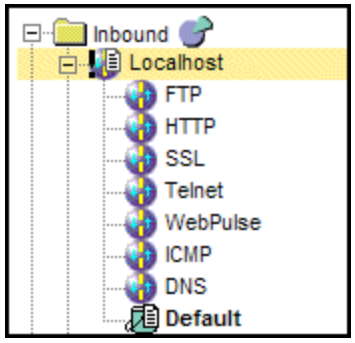

Note: If there are other Symantec products on your network that are using the WebPulse service (such as ProxySG or ProxyClient), these WebPulse queries are classified into a WebPulse class at the root (not Localhost). The Localhost/WebPulse class classifies only PacketShaper queries to WebPulse.

WebPulse Deployment Requirements

WebPulse has the following deployment requirements:

- The PacketShaper must have Internet access to connect to the WebPulse service.
- A DNS server must be configured on the PacketShaper.
- If you want to secure access to the outside interface, do not use the secure option because WebPulse requires access to a number of outside web servers. Instead, use the *list* security option and add the IP addresses of the following servers to the exception list:
    - WebPulse service points (Use the **setup webpulse show service** CLI command to see the IP addresses of the servers; add the one or two fastest servers.)
    - See "PacketShaper Port and URL Usage " on page 924 for a list of other URLs to add to the exception list.

> **Note:** To find the IP address associated with each of these servers, use the **nslookup** command (such as the **dns lookup** CLI command).

- WebPulse Express requires a PS-S500 with 64 GB of RAM and a WebPulse Express subscription.

Recommendations for Real-World Situations

Procedures to address your network and application issues

Analysis / Monitoring	Examine (not fix!) application traffic behavior
Control Strategies	Generic strategies to control application and network performance
Per-Application Strategies	Specific strategies to manage the performance of individual applications
Security	Handle security issues, such as DoS attacks and viruses
Administration	Attend to administrative issues such as scheduling

> **Note:** Managing the performance of today's traffic diversity is both a skill and an art. Each recommendation presents just one of many possible approaches. The recommendations are suggestions, and your own environment and business context heavily influence the right solution for you.

Analysis and Monitoring Recommendations

Procedures to examine your application and network behavior

Classify Network Traffic

Procedures for identifying and organizing your network traffic

The way PacketShaper organizes traffic dictates how it can analyze and control that traffic. Your traffic tree is the tool that organizes your traffic, and its configuration is a crucial choice in determining the features you have available.

For example, if your tree doesn't distinguish SAP from Oracle, then you can't measure bandwidth utilization for each. Or, if your tree doesn't distinguish traffic to your Paris office from traffic to your Oslo office, then you won't be able to compare traffic volumes or give prescribed amounts of bandwidth to each branch.

Application discovery and the creation of a corresponding traffic tree are prerequisites for understanding application behavior and controlling its performance. Consider these examples:

- A network manager denies there's any RealAudio on his network. But PacketShaper's auto-discovery could prove him wrong.
- Several departments share the cost of an application, depending on usage. But they can't analyze each department's usage unless they can first tell the difference between each department's traffic.
- All traffic to and from three remote, branch offices goes through one main site. One office tends to grab most of the WAN bandwidth, leaving the other two with poor performance. The main site could keep all branches running smoothly and equitably ... if it could recognize each office's traffic.
- Mission-critical applications running over Citrix are suffering sluggish performance due to the impact from bandwidth-hungry print traffic and file transfers. But before you can control bandwidth access differently for each application, you must first differentiate each application's traffic.

Traffic trees are very flexible. You can customize your tree's classes using many types of criteria including application, URL category, location, protocol, host lists, addresses, bit markings such as ToS, and many other strategies.

We have instructions to create a few types of traffic trees along with their pros, cons, and implications. You are not limited to these trees, but you can use them in initial strategies. Choose between an application-based, category-based, the model tree, or a location-based tree.

- Application-based tree: analyzes and controls traffic based solely on on application, protocol, or other characteristics of the traffic itself. Suitable in all circumstances *except* when managing several locations' or departments' traffic with deliberately distinct strategies. See also service-centric tree for IPv6.
- Category-based tree: monitors and controls traffic based on URL categories.
- Model tree: a pre-configured tree that organizes similar traffic into folders.
- **Location-based trees**: analyzes and controls traffic based on its source or destination *first*, then, optionally, on other criteria such as application. Suitable when managing several locations' or departments' traffic with disparate strategies.

- **Simple location-based tree**: manages only the gross bandwidth going to each location. Appropriate for occasions when your primary concern is to provision bandwidth, and you don't particularly care how it's used or how applications perform (or when performance is managed by other PacketShaper units). See also **site-based tree for IPv6 subnets**.
- **Location-based tree with per-location applications or categories** : provides the most power and flexibility in management. Takes more time to configure and usually requires a larger product model due to the number of traffic classes required.
- **Location-based tree with global applications** : manages each application with the same strategy for all locations. Most difficult tree to understand conceptually, but the easiest to configure and needs the least amount of traffic classes, matching rules, and configuration time. Provides a very scalable and easy-to-manage tree for traffic shaping but does limit visibility.

Create an Application-Based Traffic Tree

Instructions for detecting, identifying, and organizing the applications and protocols running on your network

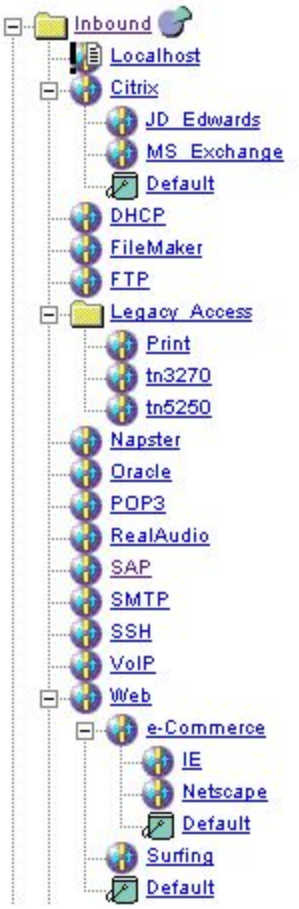

Whether you want to know which applications are running on your network, understand their behavior, track their response times, or control their performance... you always need the same first step: traffic discovery. Automatic traffic discovery and classification are a good basis for this tree's traffic classes, but you can also add or tailor classes based on other criteria. An application-based tree is a simple tree to create and use.

An application-based tree is suitable for a WAN or Internet link at any main or branch site where you want to vary traffic management strategies according to type of traffic (such as application) rather than according to destination or location. An application-based tree is most common at branch sites and main sites' Internet links.

For more information on this type of tree's capabilities, limitations, scaling considerations, and configuration recommendations, see Traffic Tree Designs.

The following steps help you create a PacketShaper traffic tree that identifies and organizes your traffic.

Steps:

1. Enable traffic discovery, telling PacketShaper to automatically classify traffic as it passes. It lets you know what's running on your network as well as distinguishes traffic so that distinct control strategies can be applied later. PacketShaper organizes the applications and protocols it detects into a hierarchical tree of traffic classes.

   For background information, see Traffic Classification Overview and Traffic Tree Overview (Advanced UI, the Sky UI).

2. Wait until enough time has passed to get a realistic representation of your network traffic. We recommend three business days if you aren't sure how long to wait.

3. Click the **Traffic Management** tab (Sky) or the **Manage** tab (Advanced UI). The list of traffic names on the left are the traffic classes PacketShaper created in response to your network traffic. **show screen...**

4. Now, or at whatever point you want your traffic tree to remain unchanged, turn traffic discovery off. PacketShaper will no longer add classes to your tree. If you want PacketShaper to continue discovering new applications within one particular class, then enable traffic discovery just for the one class (in Sky or Advanced UI).

5. Delete the classes (in Sky or Advanced UI) you don't need or want.

   Excess classes just add clutter to the traffic tree. A smaller tree will also improve PacketShaper performance.

6. Modify or add class matching rules (in Sky or Advanced UI) if you want to refine existing classes to track on additional criteria.

7. If there are multiple traffic classes that you'd like to view or manage as a group, create a folder class (in Sky or Advanced UI) and move the subordinate classes (Sky or Advanced UI) into it.

8. Create classes yourself (in Sky or Advanced UI), if discovery didn't automatically identify important applications.

9. Are there classes that you would like to define further or refine? Consult Classification Hints and Examples for ideas.

See Related Topics below for PacketGuide recommendations that describe creation of a specific type of traffic tree.

Create a Category-Based Traffic Tree

Instructions for detecting and organizing the URL categories running on your network

There are several different ways to structure a traffic tree that classifies URL categories, depending on what your goals are. Select your desired goal to read the steps for creating the corresponding tree:

- Goal 1: You want to find out which categories users are accessing.
- Goal 2: You are interested in monitoring (and perhaps blocking) traffic from specific categories.
- Goal 3: You want to treat secure and unsecure traffic for a specific category differently.
- Goal 4: You are not so much interested in monitoring/controlling *specific* categories as you are in *types* of categories.

For more information about how PacketShaper classifies URL categories, see URL Category Overview.

Create a Category-Based Traffic Tree: to Find Out Which Categories Users are Accessing

This traffic tree is built using PacketShaper's auto-discovery feature. However, unlike services, categories are not auto-discovered in the root directory of the traffic tree–you need to create a class with an "all categories" matching rule and turn on discovery for the class. As web traffic appears on your network, PacketShaper will automatically create category classes under the Categories class. Other types of traffic will be auto-discovered as service-based or port-based classes.

To create a traffic tree that auto-discoveries categories and services:

1. Decide which categories you are interested in analyzing. By default, PacketShaper discovers category classes for all categories it sees, but you can filter which categories get auto-discovered.

2. Create an "all categories" class, turn on class discovery, and make this an exception class. For details, see "Auto-Discover Category Classes" (in Sky, Advanced UI).

3. Make sure WebPulse and global traffic discovery are enabled.

4. Allow some time to pass to see what categories are discovered on your network. These categories will appear in the traffic tree as child classes of the "all categories" class.

5. View a graph of top ten categories to see which categories are using the most bandwidth.

6. Look at the Monitor screen (Advanced UI) or Traffic Management tab (Sky) to see current utilization statistics on the various categories.

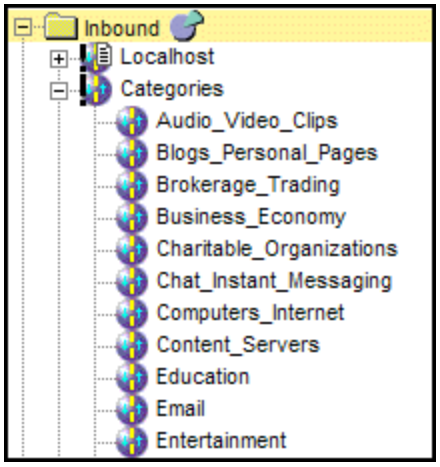

Create a Category-Based Traffic Tree: to Manage Traffic from Specific Categories

To build this type of traffic tree, you manually create classes for just the categories you want to manage. If you want to apply the same controls to multiple categories (for instance, apply a never-admit policy to all categories containing adult related content), you can create a single class (Adult_Content) and add matching rules for each category. This technique is also useful to consolidate classes in situations where your PacketShaper is reaching its maximum class limit.

To create a traffic tree to monitor and control specific URL categories:

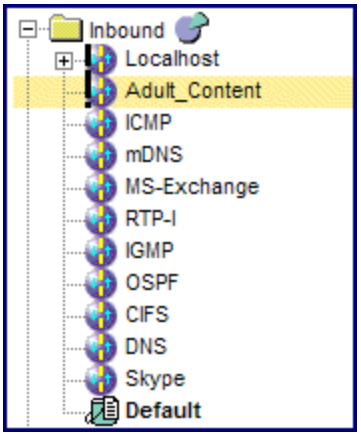

1. Create a class for the category you want to monitor.

   or

   Create a class for multiple categories.

2. (Optional) Make this an exception class.

   > **Note:** When a category class is created in the root, you will usually want to make it an exception class to make sure the web traffic gets classified into the category class instead of in service-based classes (such as HTTP, Facebook, MySpace, or other services that use HTTP as a transport).

3. (Optional) Apply a policy to the category class. See Block a URL Category if you want to block access to the websites in the categories defined in the class.

   > **Tip:** In the Sky UI, you can select multiple classes and apply the same policy to them.

4. Repeat the above steps for other category classes you want to analyze.

5. Make sure WebPulse is enabled.

6. Enable global traffic discovery if you want to auto-discover service-based and port-based classes.

> **Note:** Another way to build this type of traffic tree is to auto-discover just the categories you are interested in. By default, PacketShaper discovers category classes for all categories it sees, but you can filter which categories get auto-discovered.

Create a Category-Based Traffic Tree: to Treat Secure and Unsecure URL Categories Differently

As with the tree in Goal 1, this type of traffic tree is also built using PacketShaper's auto-discovery feature. However, instead of discovering the categories under an "all categories" class, you auto-discover categories under web-based classes. This technique allows you to have different controls for categories depending on the type of web service so that you can create different policies for categories that are secure web traffic versus unsecure. For example, you can as block unsecure financial web (HTTP) traffic, but allow secure (SSL) financial traffic. You can do this for any service that uses HTTP or HTTPS as its transport, such as HTTP, SSL, and HTTP-Tunnel.

After you have configured this type of traffic tree, PacketShaper will automatically create category classes under the web-based classes as web traffic appears on your network. Other types of traffic will be auto-discovered as service-based or port-based classes.

To create a traffic tree that auto-discovers categories under web-based classes:

1. Enable class discovery on your web-based classes.

2. (Optional) Make this an exception class.

   > **Note:** You will usually want to make this an exception class to make sure the web traffic gets classified into the this class instead of in service-based classes (such as Facebook, MySpace, or other services that use HTTP as a transport).

3. Make sure WebPulse and global traffic discovery are enabled.

4. (Optional) To apply different controls to secure and unsecure URL categories, assign the desired policies to the HTTP/Default, SSL/Default, and HTTP-Tunnel/Default classes.

   or

   Assign policies to the individual categories within each web-based class. For example, you can apply a never-admit policy to HTTP/Shopping and no policy to SSL/Shopping.

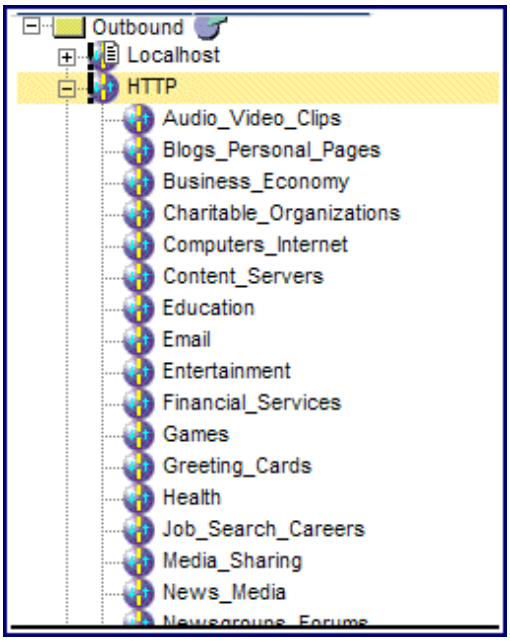

Create a Category-Based Traffic Tree: to Manage Types of URL Categories

The 70+ URL categories can be grouped into 14 category types, such as Adult Related, Social Interaction, Security Concern, Multimedia, and Leisure. For a complete list of the categories within each type, see URL Categories. If you are more concerned about controlling specific category *types* than the *individual* categories, you can create a traffic tree with category-type classes. This type of tree requires that you manually create a class with a matching rule for each URL category in the category type. For example, to create a class for the Liability Concern category type, you would add category-based matching rules for Gambling, Illegal Drugs, Illegal/Questionable, Violence/Hate/Racism, and Weapons categories. You can then control the Liability Concern categories with a single policy.

To create a traffic tree based on category types:

1. Identify which category types you are interested in monitoring and/or controlling. Refer to the URL Categories list.

2. For each category type, create a class that has a matching rule for each URL category in the type.

3. Make the class an exception class.

   > **Note:** When a category class is created in the root, you will usually want to make it an exception class to make sure the web traffic gets classified into the category class instead of in service-based classes (such as HTTP, Facebook, MySpace, or other services that use HTTP as a transport).

4. Make sure WebPulse is enabled.

5. (Optional) Assign appropriate policies to the category-type classes you want to control. To block access to websites in a category type, refer to Block a URL Category.

Use the Model Tree

The Model tree is a pre-configured traffic tree that can be used as is, or modified to suit your needs. This tree organizes network traffic into folders of typical genres, such as VoIP, risky websites, business-critical applications and data, and recreational websites and applications. It includes classes based on service groups and URL categories. You can use this tree as a starting point or an example for building your own tree. It provides an infrastructure which you can customize with your preferred configuration.

To create the Model tree:

1. Use the config save CLI command to back up your configuration before resetting the tree; this gives you the capability of restoring the traffic tree and configuration if you change your mind.

2. Reset the traffic tree, selecting the Model option.

   > **Note:** When you reset the class tree to the Model, the following global settings will take effect: Traffic Discovery will be disabled and URL Categories will be enabled.

3. To view the tree, click the **Manage** tab.

4. To understand what traffic is being classified in each part of the tree, explore the folders and classes.

   - To see all the classes in the tree, click the **Expand All** icon.
   - Click each class in the expanded tree and look at the matching rules in the right pane. The rules indicate whether the class is identifying a particular service group, multiple service groups, a particular URL category, or multiple categories. Examples:
     - In the Business-Critical folder, the Database-ERP class has a matching rule for the DatabaseERP service group.
     - For the Web-General/Commerce class, there are 11 matching rules that identify 11 URL categories (Auctions, Brokerage/Trading, and so forth).

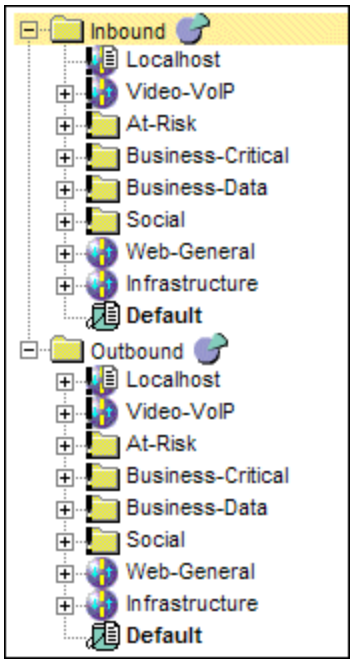

5. Some of the service group classes (such as Video-VoIP) have class discovery enabled. If you want PacketShaper to discover classes for the specific applications in that service group, you need to turn on global traffic discovery.

To customize the Model tree:

1. Go to the Monitor tab (Advanced UI) or Traffic Management tab (Sky UI) to see which classes are getting traffic hits. This lets you know the type of traffic being generated on your network. If you have global discovery enabled, you can see what classes were discovered.

2. Customize the tree by:
   - editing classes (Advanced, Sky)
   - adding classes (Advanced, Sky)
   - deleting classes (Advanced, Sky)

3. *(optional)* Based on the type and rate of traffic you are seeing on your network, set policies and partitions according to your company's objectives. See recommendations on control strategies and per-application strategies.

Create a Simple Location-Based Traffic Tree

Instructions for organizing the traffic running on your network by location only

A simple location-based tree is for a main site's WAN or Internet link with traffic that goes to multiple branches or to multiple departments. It categorizes first by travel direction and then by location.

It's appropriate for occasions when your primary concern is to provision bandwidth, and you don't particularly care how it's used or how applications perform. Additionally, it can be appropriate for a main site in topologies where other PacketShaper units manage applications at each location, and the main site just manages the amount of traffic going out to each branch.

For more information on this type of tree's capabilities, limitations, scaling considerations, and configuration recommendations, see Traffic Tree Designs.

The following steps help you create a simple location-based traffic tree.

Steps:

1. Disable traffic discovery, telling PacketShaper not to automatically classify traffic and create classes as it passes. You will create your own classes instead.

2. If you already had some application-based classes in your tree, remove them (in Sky or Advanced UI). You should have a tree with only the basic classes your PacketShaper unit had when originally shipped.

   The factory-default traffic class tree includes:

   - /Inbound and /Outbound classes with a Default class for each
   - A Localhost class for the Inbound and Outbound directions

3. Create a new class (in Sky or Advanced UI) for each of your locations under both the Inbound and Outbound branches of your tree. When you create your matching rules, use criteria to identify your locations (usually subnets, host lists, addresses, or VLAN or PVC identifiers). (Example)

   You can create each class manually. Or, if there are many classes and it gets to be a tedious task, make it easier by creating a PacketShaper command file and a data entry form.

PacketGuide recommendations that describe creation of a specific type of traffic tree:

Create a Location-Based Traffic Tree with Per-Location Applications or Categories

Instructions for organizing the traffic on your network, first by location and then by application or category

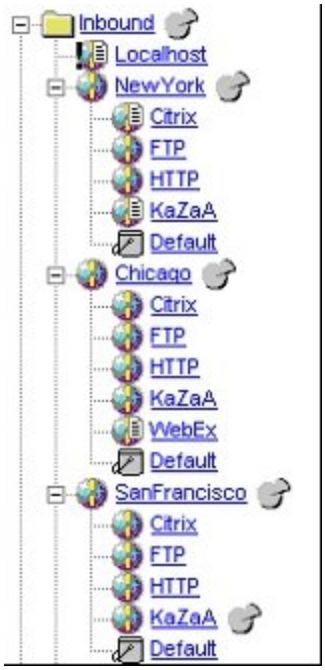

A location-based tree with per-location applications or categories is appropriate for a main site's WAN or Internet link with traffic that goes to multiple branches or departments. This tree gives you the most insight and the most control over your applications or web traffic. But it imposes scaling considerations and takes more time to configure. It categorizes first by travel direction, then by location, and finally by application or category.

For example, if you want to know how much New York traffic you have, and you want to know how much SAP traffic you have, and you want to know how much SAP traffic you have in New York then this traffic tree is your best bet.

For more information on this type of tree's capabilities, limitations, scaling considerations, and configuration recommendations, see Traffic Tree Designs.

The following steps help you create a location-based traffic tree with per-location applications or categories.

Steps:

1. Enable traffic discovery, telling PacketShaper to automatically create classes for applications and protocols as traffic passes. Then let some time elapse to give it a chance to discover the various applications that travel your network over time.

   If you are interested in seeing the categories of web traffic that are active on your network, create an "all categories" class, turn on class discovery, and make this an exception class. For details, see "Auto-Discover Category Classes" (in Sky, Advanced UI).

   Although you want a location-based tree, this step reveals the applications or categories you'll need to manage and saves time on later steps.

2. Disable traffic discovery, telling PacketShaper to stop classifying traffic.

3. Remove the traffic classes (in Sky or Advanced UI) for types of traffic you do not want to track or manage separately and for which it's okay to simply count the traffic in with a default class.

4. If you want to refine application traffic classes, consult Classification Hints and Examples for ideas.

5. Create a new class (in Sky or Advanced UI) for each of your locations under both the Inbound and Outbound branches of your tree. When you create your matching rules, use criteria to identify your locations (usually subnets, host lists, addresses, or VLAN or PVC identifiers). (Example)

   You can create each class manually. Or, if there are many classes and it gets to be a tedious task, make it easier by creating a PacketShaper command file and a data entry form.

6. Copy your group of application or category traffic classes (in Sky or Advanced UI) under each location. The bulk copy command reduces the number of steps to move each class.

   This step assumes you want to manage the same group of applications or categories under each location. If not, feel free to create and/or prune traffic classes from any individual location to customize the group of applications or categories.

7. Make sure your traffic tree has not exhausted PacketShaper's supply of traffic classes or matching rules. If that is the case, PacketShaper stops creating new classes.

   You can check estimates of system boundaries based on model. Keep in mind these are rough estimates of maximum limits that vary with real-time conditions.

   In addition, you can check your real-time system limits with the CLI command detailed below. Using the command-line interface, check your system limits. Examine the number of remaining traffic classes and matching rules to see if they are sufficient for the remaining classes you plan to incorporate. Again, these numbers are estimates based on current memory and other conditions.

   From the command line, enter: **sys limits**

You'll see results similar to these: **show screen...**

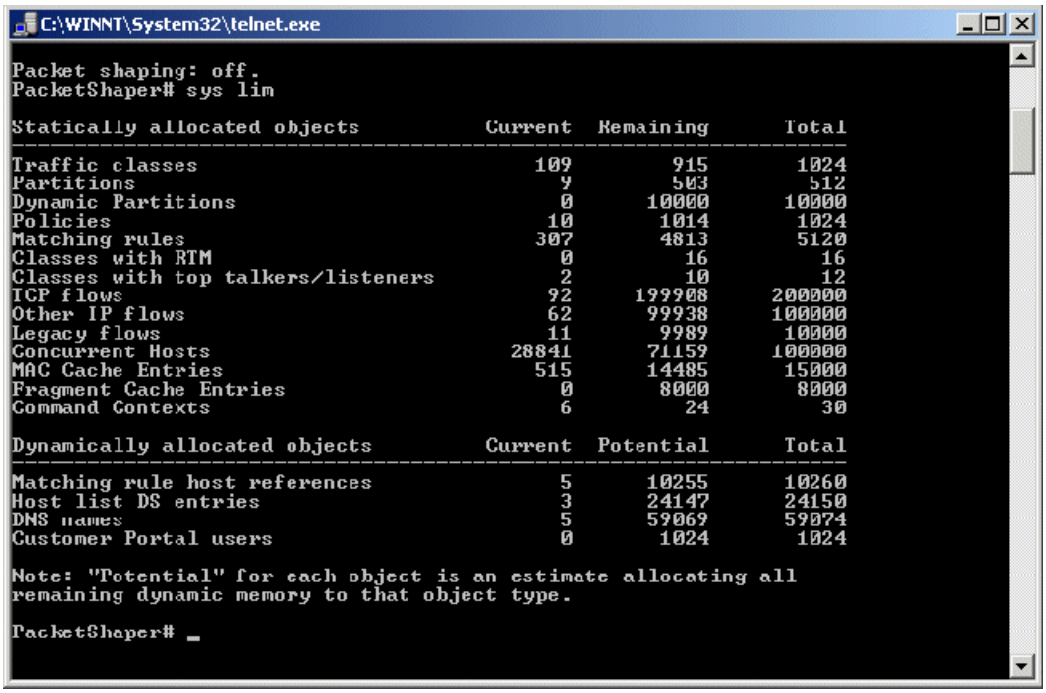

8. Remove the original group of application or category classes (in Sky or Advanced UI) that are not under any location (the group you copied to each location).

Alternatively, if you have very different applications at each location, you can create the location classes as described above, and then turn on traffic discovery for each location class (in Sky or Advanced UI). You'll get automatic traffic classes for each of your applications and protocols that PacketShaper discovers automatically. And you won't get classes for applications that are not at a particular location.

However, you will also get a **DiscoveredPorts** folder for each location with embedded port classes underneath. These can quickly consume available numbers of traffic classes and matching rules, so it is helpful to make a change to reduce the number of port classes automatically created. Normally, if the number of new connections to a static port within one minute exceeds 11, a new port class is created. Adjust the **Static Ports** variable to a larger number (200, for example) to delay the creation of each class.

When you have all the application or category classes you want, you can, if you wish, turn traffic discovery off for each location and then remove the classes you don't want. Or, you can just leave traffic discovery on to make sure you see any new applications or categores that appear. Your choice.

PacketGuide recommendations that describe creation of a specific type of traffic tree:

Create a Location-Based Traffic Tree with Global Applications

Instructions for organizing the traffic on your network by location, but applying policies by application

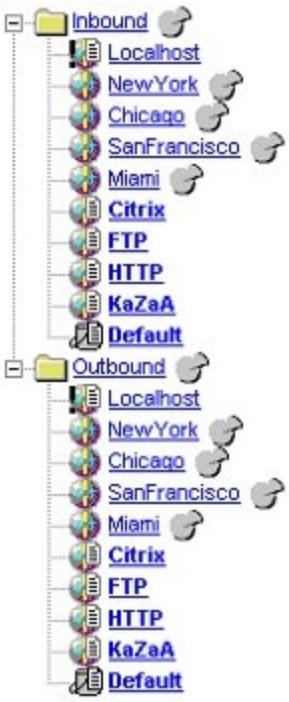

A location-based traffic tree with global applications is appropriate for a main site's WAN or Internet link with traffic that goes to multiple branches. It categorizes traffic first by travel direction, then by location. Then it applies per-application policies to each location's application traffic, borrowed from root-level inheritable classes listed below the location classes. But all traffic from one application is managed with one strategy, independent of branch.

The metrics, graphs, and reports are the same as those available in the simple location-based tree.

This tree is most common when branches have similar business priorities because each application is managed the same way for all locations. This traffic tree is the most difficult to understand conceptually, but the easiest to configure and needs the least amount of traffic classes, matching rules, and configuration time. It provides a very scalable and easy to manage main-site tree for traffic shaping but does limit visibility.

This type of tree is not appropriate for PacketShapers without the control module, as the only purpose of the application classes at the bottom is to supply policies.

For more information on this type of tree's capabilities, limitations, scaling considerations, and configuration recommendations, see Traffic Tree Designs.

The following steps help you create a location-based traffic tree with separate locations and applications.

Steps:

1. Enable traffic discovery, telling PacketShaper to automatically create classes for applications and protocols as traffic passes.

   Let some time elapse to give PacketShaper a chance to discover the various applications that travel your network over time.

2. Disable traffic discovery, telling PacketShaper to stop classifying traffic.

3. Remove the traffic classes (in Sky or Advanced UI) for types of traffic you do not want to manage separately and for which it's okay to simply count the traffic in with a default class.

> **Note:** The previous two steps are optional. You can leave traffic discovery on and refrain from pruning your classes, if you wish.

4. If you want to refine application traffic classes, consult Classification Hints and Examples for ideas.

5. Create a new class (in Sky or Advanced UI) for each of your locations under both the Inbound and Outbound branches of your tree. When you create your matching rules, use criteria to identify your locations (usually subnets, host lists, addresses, or VLAN or PVC identifiers). (Example)

   You can create each class manually. Or, if there are many classes and it gets to be a tedious task, make it easier by creating a PacketShaper command file and a data entry form.

   PacketShaper automatically places your location classes above your application classes in the tree. There is no need to make any exception classes.

> **Note:** Although your traffic tree is now complete, remember that additional steps are needed to implement your control strategy for this type of tree. The PacketGuide recommendation Control Branch Traffic from a Main Site has details.

PacketGuide recommendations that describe creation of a specific type of traffic tree:

Analyze your Link's Behavior

Instructions to evaluate high-level link behavior

Do you know which applications are using most of your bandwidth? Do you know your network's patterns of bandwidth usage? Do you know how much of your bandwidth goes to waste from retransmissions? The steps below show you techniques to help answer these questions.

Steps:

1. Look at the Top 10 pie chart (in Sky or Advanced UI) to see what traffic is currently consuming the most bandwidth on your network.

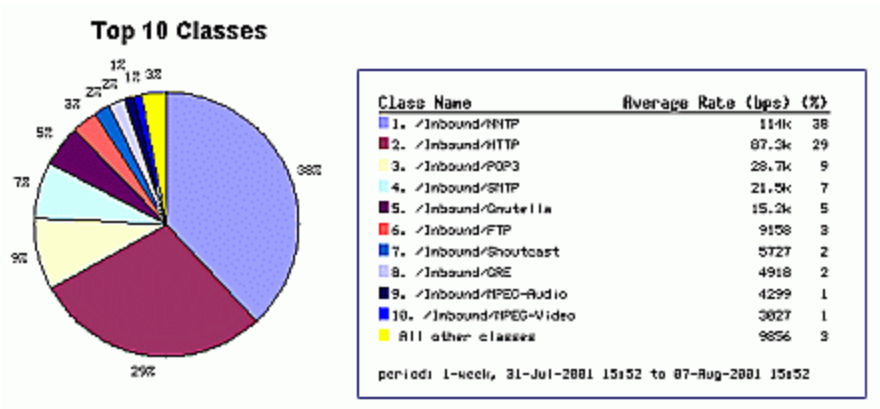

In the Top 10 chart above, NetNews (NNTP), web page downloads (HTTP), email downloads (POP3), inbound mail to your mail server (SMTP), file sharing with Gnutella, file downloads with FTP, real-time audio with Shoutcast, MPEG-Audio, and real-time video (MPEG-Video) are the most used applications.

Are there any surprises in your group of ten? Are these the applications you would expect to see as the top bandwidth consumers? Are they the group you want to see? You can use PacketShaper control features to restrict access for any traffic class whose traffic doesn't warrant the bandwidth it takes. Also, you can use control features to give bandwidth preference to critical applications that are not on your list, but should be.

2. You may want to compare the current Top 10 pie chart to the Top 10 report over the last week to make sure the current classes are representative of typical usage. To do this, display your Network Performance Summary in the Advanced UI and change the interval at the top to cover the last week. How do the lists of classes compare?

3. Now examine the Utilization graphs (in Sky or Advanced UI) for your inbound and outbound links over the last week. The graphs help you analyze how heavily your WAN link is loaded relative to its capacity, aggregating data from all applications. You can then go into a more detailed analysis of specific applications to look at their bandwidth-usage characteristics and compare them to the link's patterns.

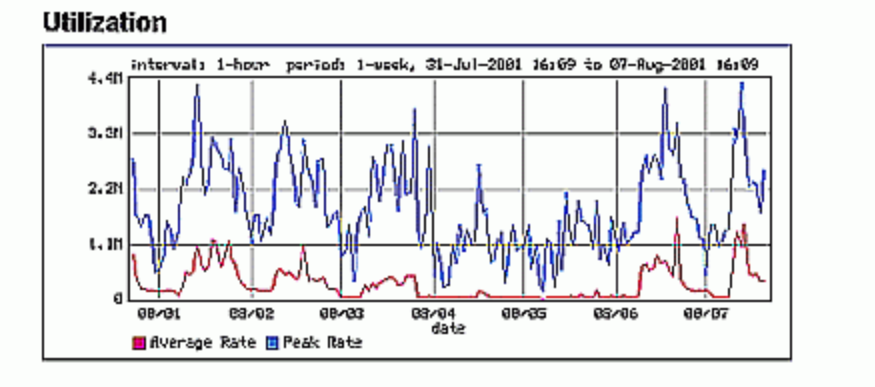

From this chart you can see that the unit has been connected to a network link with maximum capacity of about 4.5 Mbps, perhaps a network location in where three T1/E1 lines connect to a LAN.

You can also see that both average (red) and peak (blue) traffic utilization was lower during two days than the remaining five days. Perhaps it was a weekend or any other interval of lower usage.

Looking at just the red, average-rate line in the graph, note that average utilization doesn't exceed 1 Mbps on a sustained basis, even though a few peaks do go higher. Without seeing the blue, peak-rate line, this average might lead you to think that there is plenty of bandwidth to spare and congestion is never a problem. However, as you can see, peak rates frequently spike to near capacity.

Examining the last two days on the graph, more traffic are evident. Even the average line stayed up near 1 Mbps on 08/06. Perhaps one of the T1 (1.5 Mbps) lines was saturated for brief periods.

Validating these guesses would require more detailed analysis for those specific dates. When a network link has multiple lines, sometimes traffic for specific applications only travels across one of the lines. Although there is a total of 4.5 Mbps capacity, there can be cases where the maximum throughput is actually 1.5 Mbps if the traffic is not evenly shared across the three lines.

4. Display Utilization with Peaks graphs for individual classes (in Sky or Advanced UI). Investigating the traffic classes that were the top consumers in the pie chart is a good start. Adjust the graphing interval to cover one week. Consult examples of applications' bandwidth utilization graphs and their interpretations.

5. Look at the Network Efficiency graphs for Inbound and Outbound links over the last week (in Sky or Advanced UI). The graph shows the percentage of traffic that was new, first-time traffic (not retransmissions). Low efficiency means high retransmission rates, a sign of congested router queues, delays, and time-outs. (**Note:** The formula for efficiency is `bytes - tcp-retx-bytes / bytes`.)

   The chart below shows almost optimum efficiency, meaning that very little bandwidth was devoted to retransmissions.

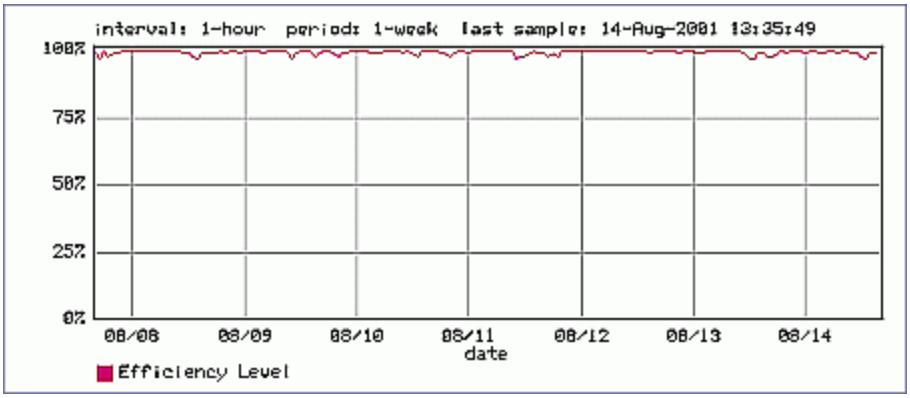

6. Display Network Efficiency graphs for individual classes (in Sky or Advanced UI). Examine the same traffic classes you investigated for bandwidth utilization, probably those that were the top consumers in your Top 10 pie chart. Adjust the graphing interval to cover one week. Consult examples of applications' efficiency graphs and their interpretations.

   For more on evaluating network efficiency, see the Assess Wasted Capacity recommendation.

Other Ideas:

If you collect a history of the Utilization graph from Network Performance Summaries over time (or the Utilization with Peaks graphs of your Inbound and Outbound links), you can estimate if and when you will outgrow your WAN capacity. Are you continually using more and more of your WAN capacity not just with peak rates, but with sustained averages, and even with appropriate policies and partitions in effect? As your usage approaches the size line, you will need to add capacity or change management strategy. For example, perhaps you've been overly generous in allowing employees unrestricted access to Internet radio, but if that means an expensive bandwidth upgrade is imminent, then you'd rather restrict access.

Assess Wasted Network Capacity

Instructions to detect, quantify, and pinpoint bandwidth wasted due to retransmissions

Retransmissions, traffic that must traverse the network multiple times for successful arrival, should optimally be as close to zero as possible. But when router queues deepen and cause dropped packets, retransmissions spike. When latency increases the frequency of time-outs, retransmissions spike. When a busy IP network behaves precisely as designed under heavy loads, retransmissions spike. Buying bandwidth to support a high rate of retransmissions is a costly, wasteful situation.

These instructions show you how to play detective with your network's retransmissions. However, you should remember that well-managed bandwidth prevents network inefficiency. PacketShaper's control features, including TCP rate control, help avoid router queues, dropped packets, and retransmissions so that you will have less inefficiency to investigate.

Steps:

1. Examine the Network Efficiency graphs for the Inbound and Outbound traffic classes by displaying the Network Performance Summary (Advanced UI) or displaying historical graphs of the Network Efficiency graph type (Sky).

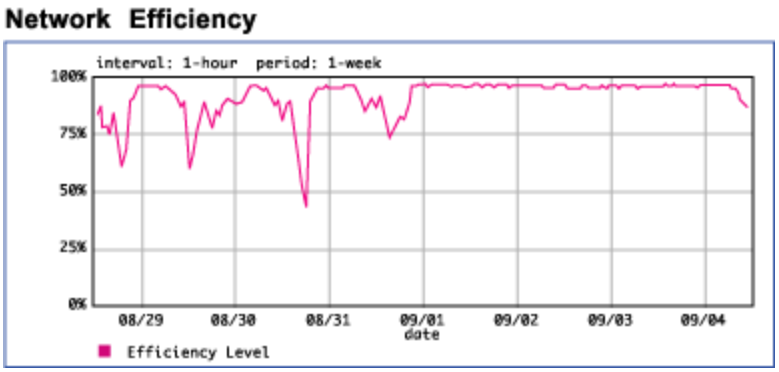

The Network Efficiency graph shows the percentage of your bandwidth wasted by retransmissions. PacketShaper knows the rates of throughput and retransmissions for any traffic class, and calculates percentages of wasted bandwidth. With the Network Efficiency graph, you can track the current retransmission rate or explore its history. You can focus on the traffic that is of interest: your link as a whole, an application, a protocol, a subnet, a user, a server, or a web destination.

Suppose you notice your Outbound class' network efficiency is fairly steady at about 90 percent.

2. Try a few recent time periods with the report's interval control settings to see if the inefficiency figure is a steady one.

For example, you might note that Outbound's efficiency figure was at 97 percent or higher until the past couple days.

3. Select several critical traffic classes and display each Network Efficiency graph (in Sky or Advanced UI). Consult examples of applications' efficiency graphs and their interpretations.

Suppose one of your efficiency graphs indicates that all classes but one show less than 4 percent retransmissions. One, the Outbound/Oracle class, shows 80 percent network efficiency or, alternatively, 20 percent retransmissions. Definitely too high.

4. At this point, some of PacketShaper's other analysis features would be appropriate. For example, you could examine Oracle's average and peak usage rates in the class tree (in Sky or Advanced UI), and analyze server and network response times for Oracle by examining the Transaction Delay graph (in Sky or Advanced UI.

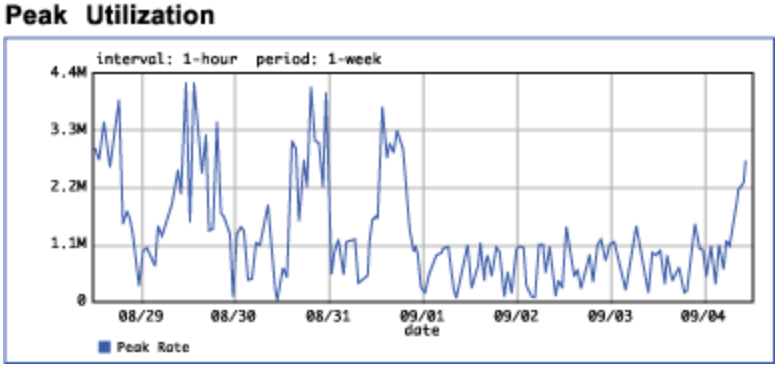

For example, you might note that Oracle's network delay has been consistent and within limits. But its server delay started climbing a couple days ago.

5. In addition to examining the Network Efficiency graph, you could display the Bytes Transmitted graph that plots the number of bytes transmitted and retransmitted together, on one graph. The result can be quite enlightening. If you have no spikes in transmissions without a parallel spike in retransmissions, then high usage is not the source of congestion problems, but error conditions (such as down or overloaded servers) are.

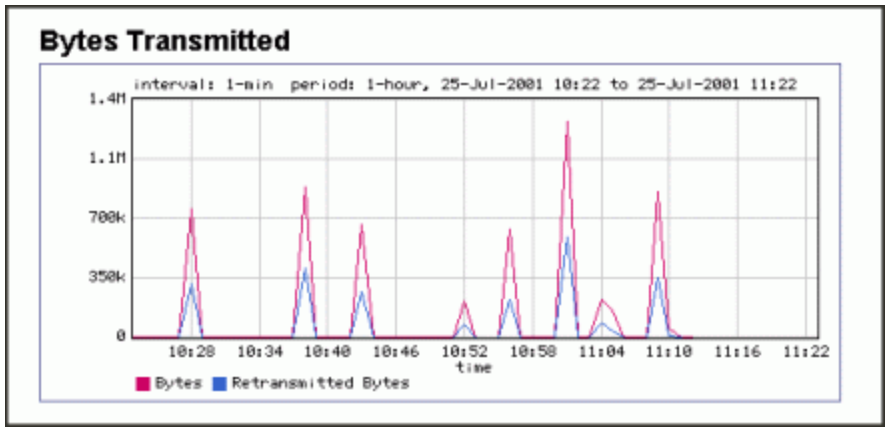

6. Choose your next investigative technique based on your current findings.

In the example, because the server delay was high, a good choice would be to enable Worst Clients and Worst Servers on the Oracle traffic class (in Advanced UI). Another option might be to divide the Outbound/Oracle class into several child classes based on client subnet, server IP address, or any other criteria that you suspect might distinguish traffic with a higher percentage retransmissions. For example, if you created child classes with different server IP addresses, you would discover which servers had higher delays (the same function as PacketShaper's Worst Servers) and higher retransmissions. For additional ideas, peruse the list of PacketShaper graphs.

Identify Performance Saboteurs

Instructions to identify traffic and applications that undermine others' performance

Some applications and protocols tend to *burst* rapidly expand to consume great quantities of bandwidth, undermining the performance of applications that are, perhaps, more urgent. This behavior doesn't necessarily imply unruly users or bad applications. The applications might be critical but tend to consume greater than an appropriate share of bandwidth (Microsoft Exchange, for example). Or they might be important, but not interactive and not particularly urgent (FTP, for example). Or they might be recreational applications (music downloads, for example).

Steps:

1. Look at the Top 10 pie chart (in Sky or Advanced UI)) to see what traffic is currently consuming the most bandwidth on your network. In the Advanced UI, you can play with the time period to get a feel for which traffic classes are consistently top bandwidth consumers.
2. Examine the usage statistics for all classes in the tree (in Sky or Advanced UI). Search for classes whose peak approaches your link capacity.

Traffic Class Name	Class Hits	Policy Hits	Current (bps)	1 Min (bps)	Peak (bps)
Inbound			780k	391k	4.3M
Localhost	10033	10033	0	0	64.0k
SameSide	10033868	10033868	0	0	1207
AFP	4	NA	0	0	60.9k
AFS	0	NA	0	0	0
Almster	54	NA	0	0	113k
AOL-IM-ICQ	2558	NA	194	91	1.3M
Audiogalaxy	17	NA	0	0	753k
EntryPoint	0	NA	0	0	0
FTP	26668	NA	212	93	3.0M
Gnutella	29695	NA	0	0	499k
Gopher	10	NA	0	0	39.2k
HTTP	11984679	NA	589k	169k	3.8M
Ident	7034	NA	0	5	18.2k
IMAP	585	NA	0	133	986k
iMesh	121	NA	0	0	87.0k
IRC	14	NA	0	0	93.7k
L2TP	2	NA	0	0	0
LDAP	9	NA	0	0	550
MPEG-Audio	233	NA	0	0	3.1M
MPEG-Video	363	NA	0	0	1.5M

In this example, note that MPEG-Audio is taking 3.1 Mbps. Unless there's a business need for it to take such a large portion of the link, it's a good candidate for PacketShaper's control features.

Although the Top 10 pie chart provides insight about which applications use the most *total* bandwidth, it doesn't highlight applications that tend to spike. An application that uses a small amount of bandwidth consistently over time can end up in the Top 10, even though it causes few performance problems for other applications. However, an application that sporadically spikes to consume the whole link for 5 or 10 minutes imposes a much more disruptive problem, even though it might not make the Top 10.

3. Display Network Efficiency and Utilization graphs for Inbound and Outbound links (in Sky or Advanced UI). Then display the same graphs for each traffic class (in Sky or Advanced UI) you identified in the previous step that had spikes near capacity.

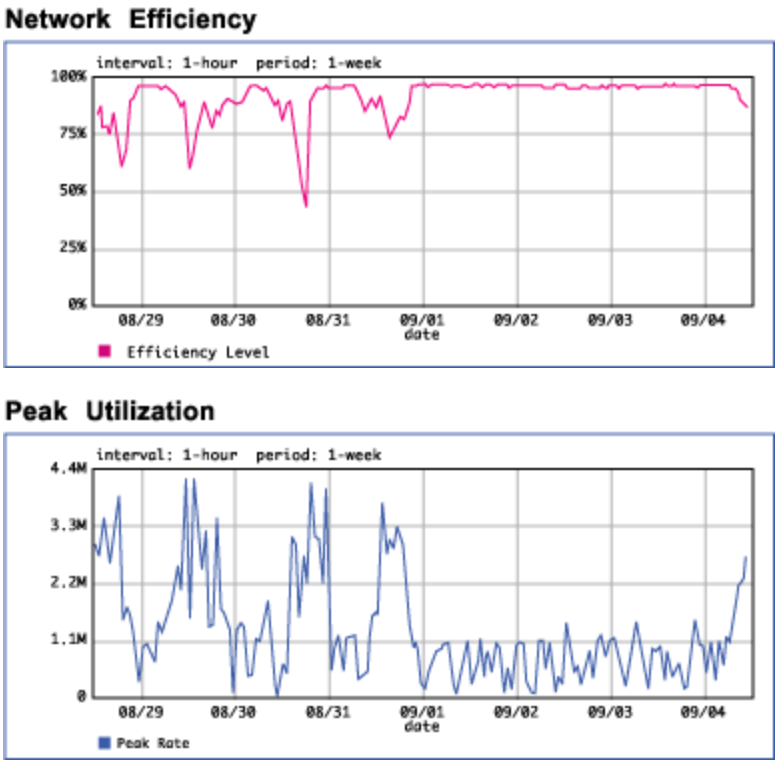

4. Compare the class graphs to the link graphs. Did class spikes impact the link? Did they bring down efficiency? Which classes caused the problems?

5. Repeat the previous step for traffic classes in your Top 10 pie chart.

   If you have enough information about users who use the most bandwidth resources, you can stop here.

6. Track the hosts that contribute most to any traffic class by setting PacketShaper's Top Talkers and/or Top Listeners features on the traffic class of interest. Top Talkers displays the top 20 generators of a class' traffic, and Top Listeners displays the top ten recipients of a class' traffic. Once set, both of them have an associated icon that appears in the Monitor Traffic window (for Top Listeners and for Top Talkers).

   A sample Top Listeners list for an HTTP class appears below.

View Top Listeners for /Outbound/HTTP — Time analyzed: 00:00:38

DNS Name	IP Address	Usage	Classify
nuq04s07-in-f27.1e100.net	74.125.224.91	41%	Add Class...
No such name	72.21.215.232	10%	Add Class...
a96-17-148-25.deploy.akamaitechnol...	96.17.148.25	10%	Add Class...
nuq04s06-in-f27.1e100.net	74.125.224.59	8%	Add Class...
No such name	138.108.6.20	7%	Add Class...
server-205-251-203-63.lax3.r.cloudfro...	205.251.203.63	6%	Add Class...

Examples of how to use Top Talkers (TT) and Top Listeners (TL):

- Use TT on HTTP to reveal the most popular websites.
- Use TL on a traffic class for the three most popular sports sites to reveal your top sports fans.
- Use TL on the current music-download application's traffic class to reveal those that retrieve the most music.
- Use TT on the same class to reveal servers most tapped as music sources.
- Use TT on an Oracle traffic class to reveal the most active Oracle servers.

Remember that PacketShaper's monitoring and analysis features are only the first half of managing performance. Control features can solve the problems that analysis only identifies.

Analyze an Application's Behavior

Instructions to evaluate how one application or one category of traffic behaves with respect to bandwidth utilization, response time, efficiency, top users, and more

Steps:

1. Create a traffic class (in Sky or Advanced UI) if one does not already exist for the traffic you want to analyze.

   For example, if you want to analyze the behavior of file transfers, then you probably don't need to do anything because the presence of FTP traffic automatically prompts PacketShaper to create FTP classes.

   But if you want to analyze file transfers from a certain server, then create a new traffic class under the FTP class that sub-classifies FTP traffic by server (for example, a class called Inbound/FTP/ServerName).

   For more information, see Traffic Tree Overview (Advanced UI, the Sky UI) and Traffic Classification Overview.

2. Make sure enough time has elapsed for PacketShaper to gather sufficient metrics for the class. If in doubt, wait for two or more business days.

3. Determine your application's peak and average bandwidth consumption over time. Use the Analyze Bandwidth Usage Patterns recommendation for help.

4. If your application is interactive, measure its performance with end-to-end response times. Or determine if slowdowns are due to slow servers or a sluggish network. Use the Analyze Application Response Times recommendation for help.

5. Determine how much of the bandwidth devoted to your application is actually traffic that is traversing the network multiple times to result in successful receipt (retransmissions). Use the Evaluate an Application's Efficiency recommendation for help.

6. Track the hosts that contribute most traffic to your traffic class by setting PacketShaper's Top Talkers and/or Top Listeners features. Top Talkers displays the top 20 generators of a class' traffic, and Top Listeners displays the top ten recipients of a class' traffic. Once set, both of them have an associated icon that appears in the Monitor Traffic window (for Top Listeners and for Top Talkers).

   A sample Top Listeners list for an HTTP class appears below.

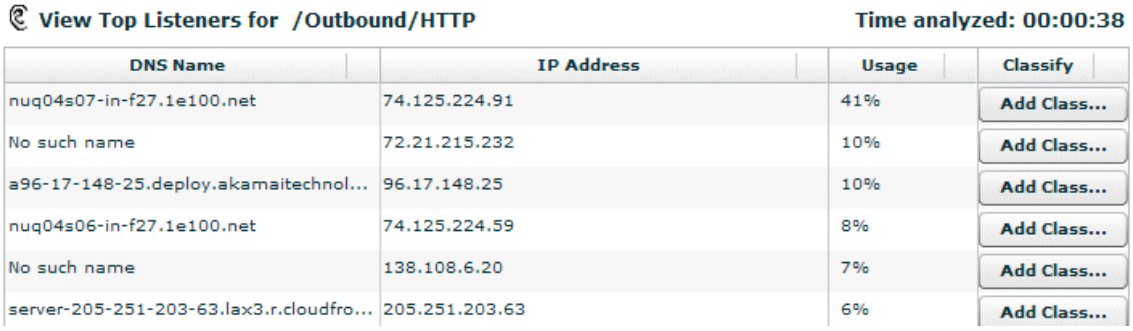
View Top Listeners for /Outbound/HTTP

Time analyzed: 00:00:38

DNS Name	IP Address	Usage	Classify
nuq04s07-in-f27.1e100.net	74.125.224.91	41%	Add Class...
No such name	72.21.215.232	10%	Add Class...
a96-17-148-25.deploy.akamaitechnol...	96.17.148.25	10%	Add Class...
nuq04s06-in-f27.1e100.net	74.125.224.59	8%	Add Class...
No such name	138.108.6.20	7%	Add Class...
server-205-251-203-63.lax3.r.cloudfro...	205.251.203.63	6%	Add Class...

Examples of how to use Top Talkers (TT) and Top Listeners (TL):

- Use TT on HTTP to reveal the most popular websites.
- Use TL on a traffic class for the three most popular sports sites to reveal your top sports fans.
- Use TL on the current music-download application's traffic class to reveal those that retrieve the most music.
- Use TT on the same class to reveal servers most tapped as music sources.
- Use TT on an Oracle traffic class to reveal the most active Oracle servers.

7. If your application is one of the ten highest bandwidth consumers on your network, use the **Top Ten** tab or the Top Ten pie chart (in Sky or Advanced UI) to analyze its usage with respect to the other nine.

8. If you want even more detail on your application than these steps have provided, try a few of PacketShaper's command-line commands. For example, explore traffic flow to see the status of all active connections for your traffic class or traffic history to see details on your application's recent flows.

Analyze Bandwidth Usage Patterns

Instructions to evaluate how bandwidth is used by an application or other category of traffic

Steps:

1. Create a traffic class (in Sky or Advanced UI) if one does not already exist for the traffic you want to analyze.

   For example, if you want to analyze usage patterns for file transfers, then you probably don't need to do anything because the presence of FTP traffic automatically prompts PacketShaper to create FTP classes.

   But if you want to analyze file transfers from a certain server, then create a new traffic class under the FTP class that sub-classifies FTP traffic by server (for example, a class called Inbound/FTP/ServerName).

   For more information, see Traffic Tree Overview (Advanced UI, the Sky UI) and Traffic Classification Overview.

2. Ensure that PacketShaper has had sufficient time to measure usage patterns. If in doubt, wait for at least a couple days. A week is even better.

3. In the Advanced UI, go to the **Report** tab to display the Network Performance Summary. Change the interval at the top to cover the last week. Find the Top 10 pie charts. Check to see if your application is listed in the highest bandwidth consumers.

   Evaluate whether you think it's appropriate that your application be listed or not. Is it a critical application that needs definite bandwidth, such as SAP, Oracle, or a Citrix application? Then it's probably appropriate to be listed. Is it a critical application that uses tiny amounts of bandwidth, such as Telnet? Then it's okay not to be listed. Is it a bandwidth-greedy but superfluous application such as music sharing? Then you'd prefer not to see it listed.

4. Create a graph (in Sky or Advanced UI) of bandwidth utilization with peaks for your traffic class. Adjust the graphing interval to cover one week.

   Do the peaks extend near link capacity? Then your class' bandwidth consumption is probably impacting other applications' performance at least occasionally.

   Is the average line much lower than the peak line? Then your traffic has sporadic spikes, but the class' typical consumption isn't so impacting.

   Does you class' usage correspond to personnel's work schedules? Are their identifiable patterns?

   Consult examples of applications' bandwidth utilization graphs and their interpretations.

5. Compare your application's utilization graph with the one for your whole link on your Network Performance Summary. Do your application's peaks impact the link's peaks?

6. Examine the statistics in the class tree (in Sky or Advanced UI). Find the row for your traffic class.

   Examine the current rate and the peak rate for your class. In this example, the Inbound class folder on the Monitor

Traffic screen shows 3.1M as a **Peak** value, which is probably two T1 lines. The class for inbound mail never uses more than 1.5 Mbps, so it probably can access only one T1 line.

Traffic Class Name	Class Hits	Policy Hits	Current (bps)	1 Min (bps)	Peak (bps)
Inbound			1.5M	256k	3.1M

Traffic Class Name	Class Hits	Policy Hits	Current (bps)	1 Min (bps)	Peak (bps)
POP3	446433	NA	10.7k	67.4k	1.5M

7. If you want to explore the top contributors and consumers of the traffic in your class, set Top Talkers and/or Top Listeners on your class (in Advanced UI only). Then wait for more traffic to pass and examine the results.

   Examples of how Top Talkers (TT) and Top Listeners (TL) can be handy:

   - Use TT on an Oracle traffic class to reveal the most active Oracle servers.
   - Use TL on a traffic class for the three most popular sports sites to reveal your top sports fans.
   - Use TL on the current music-download application's traffic class to reveal those who retrieve the most music.
   - Use TT on the same class to reveal servers most tapped as music sources.

8. You can compare bandwidth utilization in several classes by graphing multiple classes (in Sky or Advanced UI). Choose the metrics for utilization average and peak.

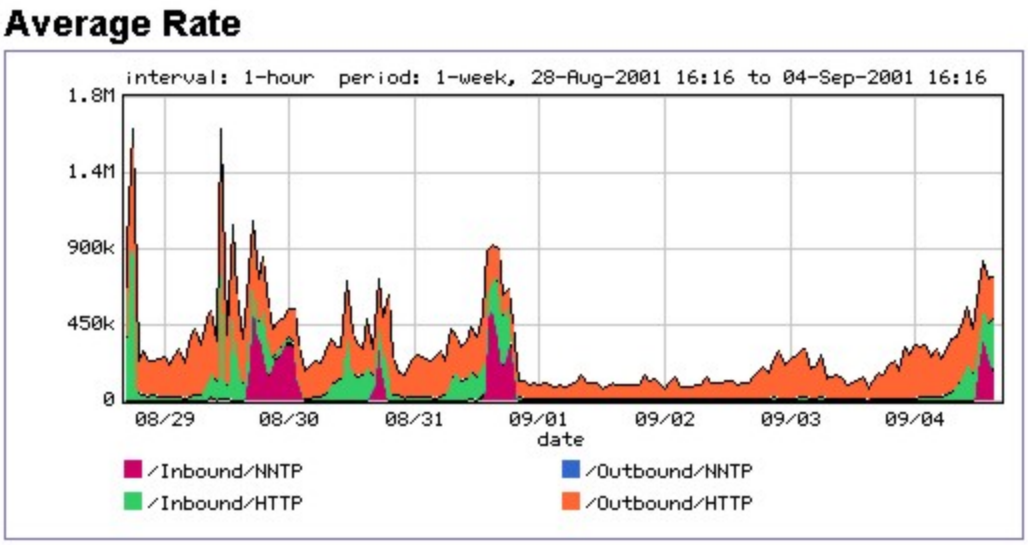

Evaluate an Application's Efficiency

Instructions to evaluate how much of the bandwidth for one application or type of traffic is due to retransmissions

Steps:

1. Create a traffic class (in Sky or Advanced UI) if one does not already exist for the traffic you want to analyze.

   For more information, see Traffic Tree Overview (Advanced UI, the Sky UI) and Traffic Classification Overview.

2. Select your traffic class and display its Network Efficiency graph (in Sky or Advanced UI).

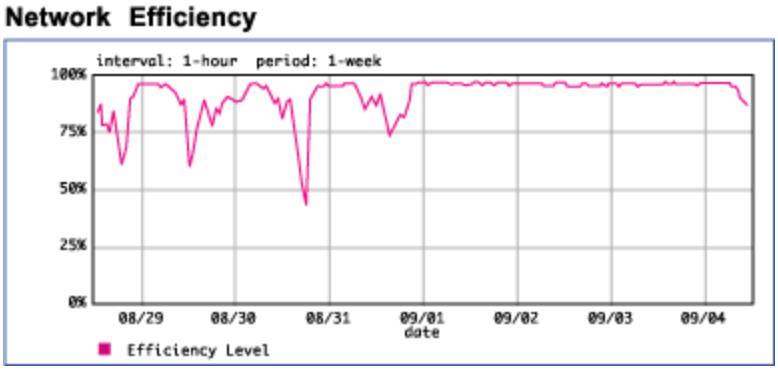

PacketShaper's Network Efficiency graph shows the percentage of bandwidth wasted by retransmissions. PacketShaper knows the rates of throughput and retransmissions for any traffic class, and calculates percentages of wasted bandwidth. With the Network Efficiency graph, you can track the current retransmission rate or explore its history.

When you create a Network Efficiency graph on one traffic class, PacketShaper reports the percentage of retransmissions for only the traffic belonging to the class.

Consult examples of applications' efficiency graphs and their interpretations to see how to interpret efficiency graphs.

Well-managed bandwidth prevents network inefficiency. PacketShaper's control features, including TCP rate control, help avoid router queues, time-outs, and dropped packets so that you have less retransmissions.

Analyze Application Response Times

Instructions to assess the network, server, and overall response times for one type of traffic

The position of the PacketShaper in the corporate network monitoring and controlling all the traffic that passes gives it a unique opportunity to provide accurate response-time measurements at a very low cost. Because it already classifies every packet, the PacketShaper can easily calculate the time traffic spends traveling between a client and a server and the time used by the server itself. Rather than collecting response data, the PacketShaper notes response times as traffic passes.

Steps:

1. Determine if you want to measure response times for:

   - ***Real traffic in real users' transactions***
     Choose a traffic class for which you'd like to investigate response times, and continue with the next step of this recommendation. For example, if you have HTTP traffic, you could analyze response times on the web. If you have an traffic class for an ERP application such as SAP, you could choose that.

   - ***Automated, manufactured, standardized, synthetic transactions***, created for the explicit purpose of measuring performance
     Initiate standardized, artificial traffic, creating corresponding traffic classes, just to analyze behavior. Then continue with this recommendation using the classes for your synthetic transactions.

2. Examine your Monitor Response Time window.

   This window lists a variety of performance-related information for each traffic class with response-time statistics.

   PacketShaper measures response times for TCP transactions. Transactions are inherently bi-directional with both inbound and outbound traffic. Although an application usually has two traffic classes (one for each direction), PacketShaper chooses one class (the one representing the client-to-server travel direction) to track and display the round-trip transaction response times.

   For example, you'll find response-time metrics in /Outbound/HTTP for external web-page transactions (outbound URL requests and their associated inbound web content). And that's the class that appears on this Monitor Response Time screen.

3. Find the row for your class of interest. Note the figures in the columns for average total delay, server delay and network delay.

   The total delay reflects all time required to service a transaction request and return the response. This time consists primarily of the time transactions spend on the network and the time the server requires to process the request. A sluggish network, an overloaded server, or a particularly large number of packets per transaction can all delay a transaction.

   Suppose, for example, your HTTP's total delay shows 3700 milliseconds. That means that your web transactions used over three and a half seconds, on average, to complete. That's pretty slow. Suppose the average network delay is 3450 milliseconds and the average server delay is 250 milliseconds. This tells you that the network is the main time-taker for

your web transactions.

4. Find the figure for your class' round-trip time (RTT). This figure represents the time for precisely one packet to travel from client to server and back. It removes the factor of transactions' size and the impact of size on network delay. If RTT is significantly smaller than the network delay, then sizable flows are to blame for inflated network delay rather than a sluggish network.

5. View your class' response-time graphs and adjust the interval at the top to cover the time of interest.

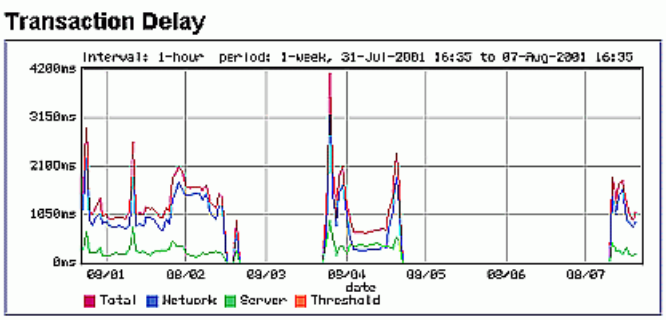

These graphs give you a visual representation of your traffic's delay figures over time and highlight the good periods and the bad. Are there specific times that correspond to poor performance?

6. If your server delay is slow, set Worst Servers on your traffic class to determine which servers are delivering the slowest processing times.

7. If your network delay is slow, find out what else is going on during these slow times. Display the utilization graphs for the entire link for the same interval of time as your Transaction Delay graph. Compare the two graphs. Do the spikes in overall utilization correspond to the spikes in network delay?

8. Reset the time period at the top of your Network Performance Summary window to cover only the interval during one of the peaks in delay time and utilization. Then examine Top Ten pie chart. Which applications were the main contributors to the link utilization during this time of slow performance? These are some of the applications that need better bandwidth management.

Slow response times can be improved. If the culprit is server delay, then you can pursue options such as load balancing and application tuning. If the culprit is network delay, then PacketShaper can help. Follow these steps:

1. Figure out which traffic is undermining the performance of your critical traffic with the techniques above or those in the Identify Performance Saboteurs recommendation.

2. Put a protective control strategy in place to throttle back truly greedy traffic, pace important traffic that shouldn't dominate, protect urgent traffic, and so.

3. Define the threshold that separates good performance from bad performance for your traffic of interest.

4. Specify the percentage of transactions that should stay within your definition of "good" performance.

5. And finally, keep tabs on your application's performance by monitoring its adherence to your performance goals. You can graph the Service Level Compliance graph.

Monitor a Condition of Interest
Notify Someone about a Condition of Interest
Respond to a Condition of Interest
... All Automatically

Procedures to have PacketShaper automatically monitor any of a large set of potential problems and then respond by notifying somebody or taking corrective actions

- Retransmissions rise to 30 percent of your network
- A single host monopolizes a substantial portion of your network
- You run out of room for additional partitions or traffic classes
- SAP's response time dips below promised service levels
- Too many users swamp a bandwidth-provisioning service, and there's no room for more

To avoid network surprises, PacketShaper automatically monitors a wide variety of predefined conditions. You can also have PacketShaper monitor your own custom conditions. PacketShaper can notify somebody and/or take more proactive corrective actions in response to satisfied conditions. For background information, see Adaptive Response Overview.

To choose the appropriate recommendation for instructions, follow these guidelines:

1. Peruse the list of adaptive response agent templates to see if one of the templates matches the condition you want to monitor. If so, follow the instructions in Monitor and Respond to Conditions of Interest using Agent Templates.

2. If you didn't find a template that suits your needs, then you'll need to define a condition based on the value of a PacketShaper measurement variable. Consult Monitor and Respond to My Own Custom Condition.

Monitor and Respond to Conditions of Interest using Agent Templates

Instructions to have PacketShaper automatically monitor any of a large set of potential problems and then optionally respond with notification or corrective actions

PacketShaper's adaptive response feature automatically monitors for conditions of interest, detects potential problems, notifies somebody if a problem is detected, and takes corrective actions. You get to decide:

- what constitutes a problem
- if you want notification and, if so, by what method
- if a corrective action is needed and, if so, which actions are appropriate

Phase One: Plan Your Adaptive Response Agent

1. Peruse the list of adaptive response agent templates to find the one that interests you. Start a page of notes for your agent plan. Record the agent name as the first item on your plan.

   There are several groups of PacketShaper-provided adaptive response agents that you can customize, including categories for the health of the network, the unit itself, and applications.

   If there is not a PacketShaper-provided agent that monitors for the condition you have in mind, you can create your own agent using any of the over 100 metrics PacketShaper measures and stores. If you want to create a metric-based agent, consult the Monitor and Respond to My Own Custom Condition recommendation instead of this one.

   We'll use the following scenario as an example throughout the rest of the steps in this recommendation:

   > Suppose you want to detect any user who uses a really excessive amount of bandwidth. You'd like to drastically limit their access to network resources, at least for the rest of the day. You decide that although you'd like all users to stay below 5 percent of capacity, you won't take action unless someone exceeds 12 percent.
   >
   > For our example, you notice that the agent template called High Bandwidth Host looks useful, and you jot its name on your plan.

2. Check to see if you can have multiple instances of the agent that interests you. This information is in the same agent template list you used in the previous step.

   You can have only one instance of some types of agents, and you can have multiple instances of others. For example, the agent that checks for bandwidth-greedy hosts can have only one instance. That makes sense because you certainly don't need to check for greedy users twice at the same time. But the agent that checks on retransmissions can have multiple instances. This also makes sense because you can check one, some, or many different traffic classes for individual levels of retransmissions.

   If there can be multiple instances, and there is not already an agent monitoring the target you have in mind for your condition, then you will *add* an agent. Otherwise, you will *edit* an agent. Make a note on your plan as to whether you will add or edit your agent.

3. Explore the information about your chosen agent from the same agent template list you used in previous steps. Then determine the agent thresholds and parameters you'd like to use.

   What condition indicates a problem or a need to respond (the red threshold)? What should indicate that things are okay (the green threshold)? Are the agent's default values for the red and green thresholds reasonable for your needs?

   For our example, you want to catch users whose usage exceeds 12 percent of capacity, and the default red value is 10. So, you note that you'll need to change that. You decide that you'd like users to stay below 5 percent of capacity, and the default green value is indeed 5. No change will be needed there.

4. What actions do you want to take when your problem gets bad (crosses the red threshold)?

   Do you want to notify somebody? Do you want to execute some other commands to document or even correct the problem?

   An *action file* is essentially a script that contains the actions you decide are appropriate. An action file can be any sequence of PacketShaper CLI commands. PacketShaper automatically executes the commands in your action file when your threshold is crossed. You don't necessarily need to take any action. Many times, just being able to check if your condition transpired is enough.

   Each agent has a set of *variables* whose values are passed along at run time to any triggered action file. The action file can use any of those variables' values, or not (your choice). You might think of the variables as parameters to the action file to avoid confusion with PacketShaper measurement variables. Consult the Action Files Overview to see which variables (or parameters) your agent makes available.

   For our example, if a host exceeds 12 percent of capacity, you want to cap their access in the future. That will be done by moving their traffic into a traffic class with per-user limits. But for our purposes, the task you need to enforce when a user gets greedy is to add their IP address to an existing host list called Abusers.

   Continuing our example, you note that $host-ip is one of High Bandwidth Host's variables. So your red action will indeed have access to the greedy host's IP address in order to add it to the host list.

5. Decide on an appropriate action or notification (if any) for when your problem returns to normal levels and stops being an issue. Again, the action file can be any sequence of PacketShaper CLI commands and can use any of the variable values made available by the agent.

   For our example, it would seem that when a formerly greedy bandwidth abuser's usage declines to 5 percent of capacity, you would want to stop restricting their access. But, that is probably not a great idea.

   You are limiting your greedy user's access, so of course usage would drop dramatically. If you restore full access as soon as usage drops, your abuser would be liberated almost immediately, and that wouldn't make for a very effective deterrent. So you will have to attend to the task of freeing greedy users somewhere else (other than in a green action).

   Make sure you think through your own situation's planned actions with the same amount of scrutiny we did with the sample scenario. Automatic responses to detected conditions are powerful, but use their power wisely to avoid any potential snafus.

6. How often would you like PacketShaper to automatically check for your condition, called an *evaluation interval*? Check the agent's default value in the same agent template list you used in previous steps. Make a note of the interval you want.

   For our example, the default interval is just a minute, and that's fine. But if, for example, it were one week instead, then that might be a problem. A greedy user could wreak a lot of havoc and impact many others in a week's time before being caught.

7. Make a note of any other details or tasks to which you must attend, if any, to complete your plan.

   For our example, you need to figure out how you will restore users to the network's good graces. Because you decided not to do so in the green action, where should it happen? You have a couple choices.

   You could add another command to the red action file. It could schedule (using the CLI **schedule** command) the *future* removal of the greedy host from the Abusers host list. Alternatively, you could automatically clear the Abusers host list at the end of each day. We'll choose the first alternative here because it uses an adaptive response feature, the focus of this recommendation.

8. With your notes intact, proceed to the next phase.

Phase Two: Create your Action Files

You can create your agent and action files in either order. You'll create your action files next here because it's convenient to have your action file names when you define your agent.

1. You can use three styles of notification (email, SNMP trap, syslog) in your action files. Each requires some setup before the corresponding CLI command will succeed. If you chose any notification to be included in either action file, attend to the required setup, as follows:

   - If you want to be notified via SNMP traps, configure PacketShaper for SNMP support and define the SNMP trap destination(s).
   - If you want to be notified via email, set up email notification by configuring an SMTP server and defining email recipients.
   - If you want a message to be sent to a syslog server, set up the syslog feature .

2. If your plan calls for a red action, you must create the action file that will be executed automatically when your red threshold is crossed.

   For our example, you would do the following:

   - Using the Info tab's File Browser, navigate to 9.258/Agent/Cmd.
   - Click **new cmd file** and enter a name in the File Name field: BADHOST.

- Enter your script in the Contents area of the Command File Editor. It would look like this:
  #Title: Restrict a greedy host by moving it to the Abusers host list (and remove it later)
  hl add abusers $host-IP
  schedule new today 23:55 "hl delete abusers $host-IP"

3. If your plan calls for a green action, you must create the action file that will be executed automatically when your green threshold is crossed (from a yellow or red status).

   For our example, you would not create a green action file.

4. Test your action files with the **run** button, supplying sample values for any $variables you used in your script of CLI commands. PacketShaper automatically prompts you for sample values and, once entered, displays the output in the results area of the page.

5. Note the names and spellings of your two action files on your plan notes. You will need them later when modifying your agent.

6. If you want to use the same action file in multiple agents, but are stumped because you would need different variables/parameters in the script's CLI commands on different occasions, you can still pursue your plan. You will need to define your own parameters in your action file. Basically, you define a parameter by using its name preceded with a **$** in your action script. Then, you supply a value when you define your agent.

   For example, suppose you want to create one action file that sends an email. You'd like to use the same action file for at least five different agents. But you want to send the email to different recipients, depending on the agent. The application manager needs the email from agents monitoring applications, and the network manager needs the email from agents monitoring network functions. You'd enter the CLI command to send email:
   send email $emailaddress $subject_heading $message_to_admin
   and, in doing so, define three new user-defined variables. Later, when you edit your agents that reference your action file, you'll supply the parameter values for any user-defined variables.

   The Action Files Overview has explanations and examples of user-defined variables/parameters.

Phase Three: Add or Edit your Agent

1. Consult your plan to see if you need to *add* or *edit* your agent. Accordingly, add an agent or edit an agent.

   Supply a name for your new agent, if you are creating a new one.

2. Fill in the agent's parameters with values from your plan notes.

   Revise the red or green thresholds, if the values in your plan are different than the values already there. If you defined red and green action files, enter the filenames in the appropriate fields.

   Note: It can be prudent and less risky to first define and test your agent *without* the names of your action files (so that

actions won't be automatically triggered before you are confident of your agent). Test your agent and your action files separately, then you can add the action file names to the agent later.

Revise the evaluation interval field if the value in your plan is different than the value already in the agent.

3. If there are any parameters listed that need values, supply them.

   For example, if you are editing an agent that monitors the efficiency of a particular traffic class, then you need to enter the traffic class name. If you defined any user-defined variables in your action file (as described in the last step of the previous phase, above), you need to supply the values here.

4. Save your new agent's definition.

5. Check your agent's status occasionally and check its reports.

Monitor and Respond to My Own Custom Condition

Instructions to have PacketShaper automatically monitor a condition you define based on a PacketShaper metric, and optionally respond with notification or corrective actions

This recommendation makes use of the adaptive response feature and its ability to allow custom definition of the condition to monitor. It's for a fairly advanced user.

Used as an example throughout this recommendation, consider the following scenario.

As the application administrator for your organization, you have tried to ensure prompt performance for your group's most business-critical application SAP. Already, you have:

- Created traffic classes for inbound and outbound SAP traffic
- Assigned partitions giving SAP traffic 20 percent of your capacity
- Defined a rate policy for each SAP class with a high priority
- Defined "good performance" as an average response time of 1800 ms or better, for at least 95 percent of the time.
- Checked performance by viewing the Service-Level Compliance graph and the table of response-time statistics

Now, you'd like to be able to stop checking SAP's performance manually. Your goal is to have PacketShaper automatically monitor SAP's service-level compliance. When performance is unacceptable, you'd like to be notified via email, and you'd like the size of SAP's partition to grow to 30 percent of capacity to address the problem while you figure out what to do for the long term.

Phase One: Plan Your Adaptive Response Agent

You will be adding an adaptive response agent to monitor your custom-defined condition, one based on a PacketShaper measurement variable. First, you'll attend to planning.

1. Make sure there is not already an agent template that monitors for the condition you have in mind. If there is, use the recommendation called Monitor and Respond to a Condition of Interest Using Agent Templates instead of this one.

   For our example, there is no suitable agent template.

2. Peruse the available class, partition, and link measurement variables. Choose the metric that assesses the condition you want to monitor. Start a page of planning notes, and note your variable's name and whether it is a class, partition, or a link variable. If it is a common variable, note whether you want to monitor the status of a class, partition, or the whole link.

   For our example, you could monitor several of the response-time metrics. But for our purposes, we'll choose **service-level%**, a class variable. It represents the percentage of transactions that satisfied their performance requirement, in our case the percentage of SAP transactions that finished within 1800 ms.

3. Determine the threshold values you'd like to have in your agent.

   What condition indicates there is a problem or a need to respond? This is the agent's red threshold.

   What indicates that things are okay? This is the agent's green threshold.

   How frequently should PacketShaper check your condition? This is the agent's evaluation interval.

   For our example, you consider good performance to be when SAP transactions complete promptly at least 95 percent of the time (the green threshold). You decide that you want PacketShaper to interpret a percentage below 90 as a problem (the red threshold). For our purposes, checking every five minutes should be sufficient.

4. What actions do you want to take when your problem gets bad (crosses the red threshold)?

   Do you want to notify somebody? Do you want to execute some other commands to document or even correct the problem?

   An *action file* is essentially a script that contains the actions you decide are appropriate. An action file can be any sequence of PacketShaper CLI commands. PacketShaper automatically executes the commands in your action file when your threshold is crossed. Keep in mind that you don't necessarily need to take any action. Many times, it is sufficient just to be able to check a condition's status.

   Each agent has a set of *variables* whose values are passed along at run time to any triggered action file. The action file can use any of those variables' values, or not (your choice). You might think of the variables as parameters to the action file to avoid confusion with PacketShaper measurement variables. Consult the Action Files Overview to see which variables (or parameters) your action makes available.

   For our example, if the service-level percentage drops below 90, you want an email to yourself and you want to increase SAP partitions' sizes.

5. Decide on an appropriate action or notification (if any) for when your problem returns to normal levels and stops being an issue. Again, the action file can be any sequence of PacketShaper CLI commands and can use any of the variable values made available by the agent.

   For our example, you decide you don't really want to do anything automatically in response to a return to green status. You don't want another email because you would already be checking SAP performance if its status had been red, making sure you fixed the problem. And you plan to attend to the partitions' sizes manually the first few times SAP performance drops. You might configure an automatic action to decrease the partitions' sizes later, once you gain confidence in how things go.

6. Make a note of any other details or tasks to which you must attend, if any, to complete your plan.

7. With your plan notes intact, proceed to the next phase.

Phase Two: Create and Test your Action Files

You can create your agent and your action files in either order. You'll create your action files next here because it's convenient to have your action-file names when you define your agent.

1. You can use three styles of notification (email, SNMP trap, syslog) in your action files. Each requires some setup before the corresponding CLI command will succeed. If you chose any notification to be included in an action file, attend to the required setup, as follows:

   - If you want to be notified via SNMP traps, configure PacketShaper for SNMP support and define the SNMP trap destination(s).
   - If you want to be notified via email, set up email notification by configuring an SMTP server and defining email recipients.
   - If you want a message to be sent to a syslog server, set up the syslog feature.

2. If your plan calls for a red action, you must create the action file that will be executed automatically when your red threshold is crossed.

   For our example, you would do the following:

   - Using the Info tab's File Browser, navigate to 9.258/Agent/Cmd.
   - Click **new cmd file** and enter a name in the File Name field: RedSAP.
   - Enter your script in the Contents area of the Command File Editor. It would look like this:
     #Title: Monitor Sap's performance. Boost partition sizes when poor.
     send email <*YourEmailAddress*> "SAP Performance Dropped" "Performance for $class-id dropped to $me-value"
     partition apply inbound/sap 30% none
     partition apply outbound/sap 30% none

3. If your plan calls for a green action, create the action file that will be executed automatically when your green threshold is crossed (from a yellow or red status).

   For our example, you would not create a green action file.

4. Test your action files with the **run** button, supplying sample values for any $variables you used in your script of CLI commands. PacketShaper automatically prompts you for sample values and, once entered, displays the output in the results area of the page.

5. Note the names and spellings of your two action files on your plan notes. You will need them later when editing your agent.

Phase Three: Add Your Agent

1. Consulting your plan, add your agent.

   Select **Link ME Variables**, **Partition ME Variables**, or **Class ME Variables** for the agent template, according to your plan.
   Supply a name for your new agent that is descriptive of its purpose.
   Adjust the **Evaulation Interval** according to your plan.
   Enter the names of your red and green action files, if any.

   For our example, you would:
   Select **Class ME Variables** for the agent template.
   Enter **Monitor SAP Performance** for the agent's name.
   Change the Evaluation Interval to 5 minutes.
   Enter **RedSAP.cmd** for the red action file's name.

   Note: It can be prudent and less risky to first define and test your agent *without* the names of your action files (so that actions won't be automatically triggered before you are confident of your agent). Test your agent and your action files separately, then you can add the action file names to the agent later.

2. Click **OK and edit parms**. The *Edit Agent Entry* window opens.

3. Edit the agent with field values from your plan notes.

   Enter the red and green thresholds, the traffic class name you wish to monitor, the name of the PacketShaper measurement variable you noted, and so on, for all the fields in the window and described in Edit an Agent.

   For our example, you would fill in the fields to look like the following:

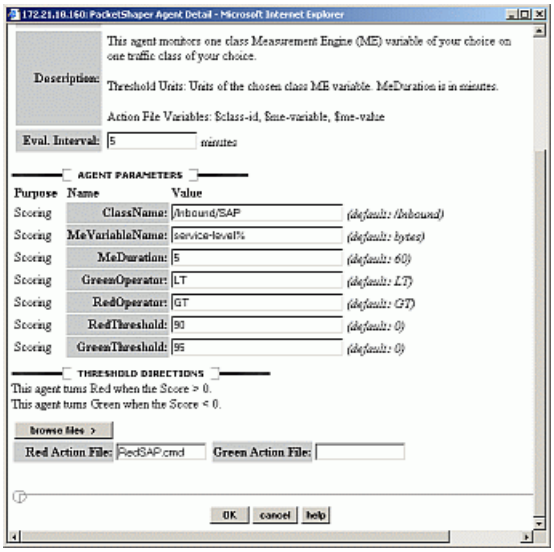

4. Save your new agent's definition.

5. Check your agent's status occasionally and check its reports.

Notify Someone of Situations of Interest

Instructions to send email, a message to a Syslog server, or an SNMP trap to indicate that a condition of interest has occurred.

PacketShaper offers several features to monitor for a condition of interest and then notify someone that the condition has occurred. This recommendation uses the older and less powerful user events feature.

Note that you must use the command-line interface (CLI) to configure events.

Steps:

1. Decide how you want to be notified when an event occurs: via SNMP traps, email, and/or syslog.

   SNMP traps are useful when using a network-management platform, such as HP OpenView, for event and alarm processing and notification. Email is an appropriate choice if you need fast, urgent notification that someone is likely to check more frequently.

2. If you want to be notified via SNMP traps, configure PacketShaper for SNMP support and define the SNMP trap destination(s).

3. If you want to be notified via email, set up email notification by configuring an SMTP server and defining email recipients.

4. If you want a message to be sent to a syslog server, set up the syslog feature.

5. Define which metrics PacketShaper will use when evaluating your situation of interest. Find a relevant predefined event or define a new event.

   For example, you could define the following event:
   event new WastedBandwidth tcp-retx-pkts%.link>$1 1m
   to define a new event called "WastedBandwidth" to be done on the link as a whole that tests the measurement variable called "tcp-retx-pkts%" for being more than something (>$1) as yet undefined. The default checking interval is one minute (1m).

   Note that you have defined only a potential event, as you still have to define the comparison value, as well as other information, and to tell PacketShaper to start monitoring.

6. Register the event, giving PacketShaper the rest of your instructions and telling it to start monitoring for your condition of interest.

   The types of information you specify include the comparison value or threshold, a specific instance of the link, partition, or class you want monitored, and criteria to help prevent a flood of notifications.

   For example, you could register the following events:
   event register WastedBandwidth(inbound,12, 4) 1m email trap limit=3
   event register WastedBandwidth(outbound,12, 4) 1m email trap limit=3
   to apply your wasted-bandwidth event to both Inbound and Outbound links, evaluate them every minute, send both an email and an SNMP trap when retransmissions grow to above 12 percent, avoid additional notification until after retransmissions decline to 4 percent, and never send more than three notifications in one day.

7. Keep tabs on all the armed events and their history by monitoring event status.

Define and Monitor a Set of Reports

Instructions to select and use a set of reports on a regular basis.

A carefully selected group of reports, monitored on a regular basis, can uncover performance problems, reveal historical trends, confirm or deny the efficacy of network or configuration changes, and more.

To choose your set of reports:

1. Peruse the list of PacketShaper graphs (in Sky or Advanced UI) to choose a set of graphs that track information that is important to your organization.

   For example, you might want to start with these three graphs for your Inbound and Outbound classes:

Graph Name in Sky	Graph Name in Advanced UI	Description
Utilization with Peaks	Link Utilization with Peaks and Size	Bar graph of the course of your link utilization
Top 10 (Dashboard)	Network Performance Summary Top Ten Classes	Pie chart of how your bandwidth is divided between your top ten applications
Network Efficiency	Network Performance Summary Network Efficiency	History of the portion of your link that goes simply to retransmissions

   If you have a location-based tree, you'll probably want to include a pie chart of the top ten branch locations -- Top Ten Classes or Top Ten Partitions.

2. If there are any graphs that you want but don't find in the list of PacketShaper graphs, then you can define your own graphs based on PacketShaper measurement variables. In addition, you can combine multiple classes into the same report and name them for future reference and use. Use a one-week interval, or any interval of your choice.

To generate and save all your graphs:

1. Generate each graph you selected (in Sky or Advanced UI) with a measurement interval of one week (or the measurement interval of your choice).

2. Generate each of your own custom-defined graphs with the same measurement interval.

3. Save all your graphs for future reference.

To analyze your graphs:

Use the Analysis section of PacketGuide's recommendations to find suggestions to assist you in analyzing your graphs. For example, you can find suggestions to monitor the condition of your WAN link using the three graphs suggested at the top of this recommendation.

Identify Mysterious Traffic

Instructions for identifying traffic that PacketShaper does not classify

Many customers view PacketShaper auto-discovery, or automatic traffic classification, as their favorite and most useful feature. Knowing the identity of traffic running over your network is a big first step in managing and controlling the performance of network applications.

When PacketShaper sees a traffic flow, it matches the flow's characteristics to those of each class in the traffic tree. If it finds a match, that class' metrics are incremented accordingly, and the flow is managed with the class' policies. If it does not find a match, but PacketShaper can identify the traffic, and traffic discovery is enabled, PacketShaper creates a corresponding new traffic class. There are several reasons PacketShaper might not create a class:

- Traffic discovery is disabled
- PacketShaper can't identify the traffic
- There are so many classes in the traffic tree that more are not possible
- Insufficient number of flows have passed to prompt PacketShaper to make a new class (anywhere from one to 11 flows must pass before a traffic type gets its own class, depending on the type of traffic)
- The PacketShaper unit was recently plugged in and started monitoring long sessions that were already in progress. Until a new flow starts, this traffic all counts as Default. In these cases, the solution is just to wait another few hours until you examine the tree.

In any of these cases, PacketShaper matches the flow with most appropriate **Default** class (usually Inbound's or Outbound's Default class).

Determining what traffic is in the Default class is usually not a concern or priority. But if the amount of traffic in a Default class increases precipitously, or if most of your traffic is classified in the Default class, you'll probably want to figure out what that traffic is. Is there a new music download application that is swamping an inappropriate amount of your bandwidth? Or is there a new critical application that you aren't protecting adequately? Or what?

Some of the same techniques used to identify the Default class' traffic can also be used to try to identify the traffic that PacketShaper deposits in traffic classes with port numbers as names under the **DiscoveredPorts** folder.

Steps:

1. Make sure that PacketShaper would be classifying your mysterious traffic if it could.

    - Ensure traffic discovery is enabled.

      Perhaps you don't want PacketShaper to create classes automatically because you are concerned that you'd clutter up your carefully crafted traffic tree in the process of getting the one class that interests you. If this is the case, then leave discovery off and follow the suggestions below. Once you identify the traffic, you can create a class manually.

- Make sure your traffic tree has not exhausted PacketShaper's supply of traffic classes or matching rules. If that is the case, PacketShaper ceases creating new classes.

  You can check estimates of system boundaries based on model. Keep in mind these are rough estimates of maximum limits that vary with real-time conditions. In addition, you can check your real-time system limits with the CLI command detailed below. Using the command-line interface, check your system limits. Examine the number of remaining traffic classes and matching rules to see if they are approaching zero. From the command line, enter: **sys limits**

  You'll see results similar to these: **show screen...**

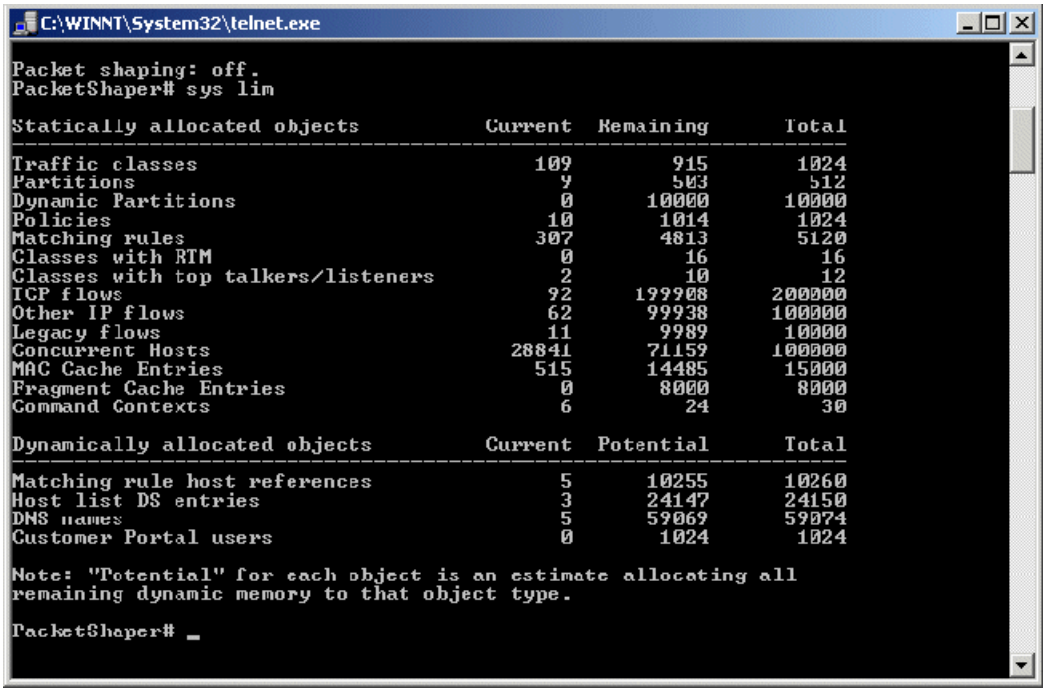

2. Set Top Talkers and Top Listeners on the mysterious default or port-based class. PacketShaper will track who are the top contributors and recipients of the traffic in the class. Let some time pass so that new traffic has a chance to pass and PacketShaper has the opportunity to monitor.

3. Examine your class' top contributors and recipients.

   If you configured a DNS server during installation and setup, then your Top Talkers and Top Listeners tables may contain domain names instead of the less decipherable IP addresses. But either way, do you recognize the clients or servers that are the top contributors? Do you know their associated applications/tasks? Can you call them on the telephone and ask them what they're doing?

4. If there is a particular contributor or recipient that is of interest, create a distinct traffic class just for its traffic. Set Top Talkers and Top Listeners on that class as well to see if the information sheds any light.

5. Use the CLI command **traffic history recent *<classname>*** for your class to see the date, time, IP address, port number, and URL for each flow in the specified class.

If there is an IP address or DNS name that you'd like to explore further, do a **traffic history find *<host>***. This command allows you to see in which classes a host's flows hit, as well as the number of flows and the protocol. If the protocol cannot be identified, a dash ( ) appears.

6. Explore the mysterious default or port-based class further with the **traffic flow** CLI command. It has many options that display a variety of different information. If you enter **traffic flow** without any any additional parameters, you get a list of the command's options.

   You can use the **traffic flow** command for a particular address that interests you (that you got from Top Listeners, for example), a particular class or host, a certain number of flows, and more. Here are some examples of useful variations on the traffic flow command:
   For a summary: **traffic flow -tulo**
   For data on 100 flows from a single class such as Inbound's Default:
   traffic flow -tupln 100 -c /inbound/default
   For a single IP address: **traffic flow -tupvA 192.168.1.10**

   Do most users communicate with a a common server or subnet? Note its IP address. Is there an associated service listed? If PacketShaper identifies the service associated with a flow (even though it must classify it in a default class), it will display the application or protocol name. In these cases, your quest ends here.

7. Repeat the same traffic flow command, this time adding the -A option for an IP address with the address you noted in the previous step. You'll filter the information further so you're not confronted with such an overload of data.

8. Take information (such as IP addresses associated with the mysterious traffic) you isolated with the **traffic flow** command to use them with industry-standard networking utilities. Suggestions include:

   - Try an **nslookup** command from a PC command prompt to get the DNS name. You can also use the PacketShaper CLI command **dns rlookup *<IPaddress>***.
   - Try **traceroute** on a PC or UNIX machine to fill out your information on the mysterious traffic's path from source to destination.
   - Find out who owns the block of IP addresses by entering the IP address at a Whois site such as http://www.arin.net/whois/index.html.
   - Explore a search site (such as Google) to search for firewall-related postings for specific port numbers.

9. As a final suggestion for more information, have PacketShaper capture a log of the class' traffic and then feed that log to a sniffer or third party analyzing software such as EtherPeek. For instructions, see the Sniff without a Sniffer recommendation.

Identify New Traffic

Instructions for quickly identifying new traffic

So that you can keep track of new applications the employees are using on your network, you can create a "bucket" into which any newly discovered classes get placed. You can periodically check this bucket to see what's new on your network, analyze how much bandwidth each new application is consuming, and then make a decision on whether you want to place controls on this traffic.

Steps:

1. Turn off class discovery for the Inbound and Outbound folder classes. (You want new traffic to get discovered in your NewTraffic match-all class.)

2. Create a class in the Inbound direction called NewTraffic that matches all IP traffic. (Service=any, Protocol Family=IP)

3. Turn on class discovery for the Inbound/NewTraffic class.

4. Create an Outbound/NewTraffic class, and enable class discovery.

5. Ensure global traffic discovery is enabled.

   New applications will now get discovered as child classes of the Inbound/NewTraffic and Outbound/NewTraffic classes. **show screen...**

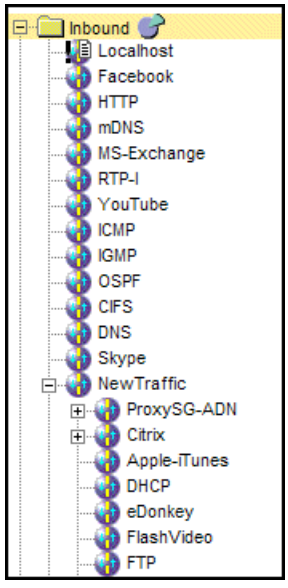

6. Review the contents of NewTraffic on a regular basis to see what's new on your network.

7. When you find an application that you want to monitor and/or control, you can move the class (in Sky or Advanced UI) out of NewTraffic into the root.

Identify Users Who Access a Category

Instructions to identify the user names or IP addresses of clients accessing a URL category

You can use the Top Talker/Listener feature to determine which users are accessing websites in non-business categories.

Steps:

1. Create a class (in Sky or Advanced UI) for the category you want to analyze.

   or

   Create a class that classifies multiple categories (Sky or Advanced UI) of similar content.

2. Track the hosts that contribute most traffic to your category class by setting PacketShaper's Top Talkers and/or Top Listeners features. Top Talkers displays the top 20 generators of a class' traffic, and Top Listeners displays the top ten recipients of a class' traffic. Once set, both of them have an associated icon that appears in the Monitor Traffic window (for Top Listeners and for Top Talkers).

3. Display the list of Top Talkers or Top Listeners for the category class.

   The Top Listeners report for an Inbound category class lists the user names (if the user awareness feature is configured) or IP addresses of the clients who accessed websites in the category. These are the hosts that have used the highest percentage of bandwidth in the class since tracking was enabled.

Sniff Without a Sniffer

Instructions for using PacketShapers to get the information typically gathered by a protocol analyzer.

Protocol analyzers are used for a variety of diagnostic purposes as they capture and analyze each passing packet. They are useful, but can be pricey if you want their benefits at all network locations. If you already have PacketShapers deployed at these locations, you need not also invest in additional analyzers at every site.

PacketShaper provides a platform that can be used to capture packets at strategic points on a network. In addition, you can even choose which traffic is logged and which is not with flexible criteria.

Many of the statistics presented by an analyzer are also available in PacketShaper pages and graphs. But if you need information that PacketShaper does not provide, you can use PacketShaper to capture all or some passing packets. Then you can read the log file with third-party analyzing software such as EtherPeek, Ethereal, or a Sniffer.

Steps:

1. Determine what type of information you want to collect.

- If you want information such as active applications and protocols, active IP addresses, bandwidth utilization, retransmissions, and heavy users, this information is available in PacketShaper. Check out the other recommendations under Analysis / Monitoring, the list of PacketShaper graphs, the Monitor Traffic window, the top hosts feature, as well as several CLI commands such as traffic flow and traffic history.

- If you want to view packet headers, real-time display of top users, content at specific offsets into packets, or other information not available from PacketShaper, you'll use the packet capture facility.

  First, you'll configure PacketShaper to capture passing packets, and then you'll pass the resulting log file to a third-party analysis tool. Continue to the next steps.

2. Decide which packets you would like to collect.

   A major advantage of using PacketShaper as a collector is that you define precisely which traffic to capture. You don't have to collect huge log files with mostly irrelevant traffic. For example, if you want to capture all Telnet packets to or from a certain IP address you can. Or if you want to capture only Oracle traffic for one particular database you can.

   You can capture packets for NICs, traffic classes, IPv4 addresses and ranges, IPv4 subnets, host lists, port numbers and ranges.

3. If you want to capture traffic for a specific class and the class doesn't already exist, create the traffic class. The class can be based on any protocol, including IPv6.

4. Add each of the NICs, classes, IPv4 addresses/ranges, IPv4 subnets, host lists, or port numbers/ranges for the traffic you want captured, one at a time, to PacketShaper's capture list with the CLI command packetcapture add. Note that packet capture has not started yet. You're just specifying the traffic that will be logged when packet capture does start.

5. If the information you want is at the beginning (or at least *not* at the end) of each traffic flow, consider limiting the number of packets that PacketShaper captures for each flow with the CLI command packetcapture limit packets. Your logs won't fill as fast, and you'll still have the information you need.

6. Turn packet capture on.

   PacketShaper stores captured packets in RAM. They are written to disk when the memory buffer is full or when you turn packet capture off.

7. If you want to monitor of the progress of your packet capture, use the packetcapture status command.

8. When enough traffic has passed, and you have enough captured data, turn off packet capture.

   PacketShaper writes the log to disk in tcpdump format and puts it in the 9.258/pktlog folder.

9. Download the log file to the computer running third-party network analyzer software such as EtherPeek or Ethereal.

10. Open your log file with your analyzer software.

Use Standardized, Artificial Traffic for Analysis

Instructions to automate synthetic transactions and collect data about them for analysis

Analyzing real-life traffic, with all its variations and eccentricities, has its strong points. It gives you insight into what really happens with your users' actual performance. But it can also introduce unpredictable results because real life is exactly that -- unpredictable.

For example, the results of measuring an application's transaction response time vary with the size of transactions, whether any users have used the application recently, and so on.

You can avoid some of these issues by automatically generating uniform transactions at regular intervals (called *synthetic transactions*) and measuring results. PacketShaper can, at your discretion, initiate synthetic web or other TCP transactions at periodic intervals to verify the availability of critical hosts as well as collect performance data.

Synthetic transactions are similar to scheduled pings or SNMP polls, but with these important advantages:

- PacketShaper's detailed analysis of transaction behavior and response times apply to synthetic transactions, giving you the ability to profile network and host behavior over time. This information is much more helpful than knowing if a device simply responds to pings or not.

- When PacketShaper sits at the network edge, polls are local, consume less bandwidth, and can therefore be more frequent and impose less network burden.

- Synthetic transactions can determine if a service or application is running, not just if a server is responding. They provide a more sophisticated assessment of application availability.

- Distributed PacketShaper appliances can serve as local clearinghouses for availability information, forwarding situations of interest or concern to central locations via email, SNMP traps, or syslog messages. The need for long-distance polling by central management platforms is eliminated.

To choose this approach, follow the steps below:

Steps:

1. To tell PacketShaper to initiate web or other types of TCP transactions at periodic intervals, access the command-line interface (CLI) and enter the command synthetic add.

   Examine and choose carefully for the many parameters in the synthetic add command. For example, suppose you want to measure the performance of web transactions to one of your Intranet application servers. You determine that a transaction once every ten minutes is sufficient, giving you 144 occurrences each day to measure, average, and assess.

   For this example, you'd choose an *interval* of 10, ignore the *repeat* and *id* parameters, choose *http*, enter the domain name or IP address of the web application server, and finish with a forward slash and an html pathname on the server. Your command would end up looking like:

synthetic add 10 "http://appserver.domainname.com/testfile.htm"

Or, in another example, suppose you want to know if a host is there and responsive. You could choose a five-minute interval, *echo,* and enter the domain name or IP address of the host you're wanting to check, rendering:

synthetic add 5 "echo://test.domain.com"

2. To analyze your synthetic transactions separately from other traffic, they must have their own traffic classes. Enter the CLI command synthetic options create-classes on to create distinct traffic classes. It's the default setting, but just in case anybody has changed the setting, it's best to make sure.

> **Note:** PacketShaper creates one synthetic transaction class per host. If you create more than one synthetic transaction for a host, all transactions will get classified in a single class. For instance, suppose you create an echo transaction to test.com and an http transaction to test.com, PacketShaper will create a single class, *Localhost/test.com*, and classify the echo and http transactions in this class.

3. Wait through a period of time that is at least a couple of your *interval* periods. Then check your **Monitor** screen.

   You should see a traffic class for your synthetic transactions with a child class for each of the hosts you included in your commands. Your host classes should list some class hits if sufficient time has passed to issue your synthetic transactions.

4. Enter the CLI command synthetic show. You'll see each of your synthetic transactions, and if any have completed, their results. For example, if you specified a synthetic transaction using *echo*, you can see if your device was responsive.

   If you decide you don't want one of your synthetic transactions, note its ID number from the output of the synthetic show CLI command, and use that ID with the synthetic delete command.

5. Decide if you want a policy on your new traffic classes for your synthetic transactions.

   You would choose a policy if you are trying to match the performance of traffic in another traffic class. If you have a class that handles your real-life transactions that are similar to your synthetic transactions, then duplicate any policy you have on the real-life traffic class.

   For example, recall the earlier scenario to measure the performance of web transactions to a certain application server. Suppose that traffic's real transactions have a rate policy (with settings of 0 bps guaranteed, burstable at priority 5, no limit). In that case, you'd want to assign the same rate policy to the test transactions under your synthetic transactions' traffic class.

   You would not bother with a policy if you are simply wanting the transactions to go through, and you don't care about mimicking any other traffic's performance conditions. For example, recall the earlier scenario to test if a device was present and responsive. In this case, you'd probably skip a policy.

   If you decided a policy was appropriate, define and apply that policy. For instructions on creating each type of policy,

choose a topic from the navigation bar in PacketGuide's Tasks->Policies page.

6. Let some time pass, at least a couple of days if you want a realistic quantity of data to analyze.

7. You can now apply any PacketShaper analysis technique to your synthetic transactions' traffic classes. Analyze response times, Check the monitor page or create any graph you see on the list of available graphs.

Control Strategy Recommendations

Procedures to control your application and network traffic

Protect Critical Application Performance

Procedures to ensure prompt, appropriate, and predictable performance for critical applications

Today's networks facilitate the simultaneous delivery of both mission-critical and unsanctioned applications. When low priority applications cause performance of critical applications to suffer, the application-neutral nature of IP networks and the sheer volume of traffic generated by many Internet-enabled applications combine to form one very infuriating performance problem.

The following procedure contains instructions to differentiate and protect your urgent traffic. Because part of that process is preventing others from taking the bandwidth that should rightfully go to the more urgent applications, it also contains instructions to limit bandwidth where needed.

Procedure:

1. Develop a traffic tree, if you don't already have one.

2. Determine which applications are undermining the performance of others.

3. Identify the applications whose performance is crucial to your organization. They are probably your interactive, latency-sensitive, mission-critical applications. Examples might include SAP, Oracle, Citrix-based applications, or a particular web page. Make sure you have traffic classes for each. If PacketShaper doesn't auto-discover your particular application, create a traffic class to identify the application with other criteria such as server address, port numbers, URL, or others.

   For more information, see Create a Traffic Class (in Sky or Advanced UI) and Matching Rules.

4. Advanced UI: Click the **Monitor** tab and sort the Monitor Traffic list by peak bps. Any class that has peaked to more than 50 percent of the link size should be assigned a policy. Make a list.

   the Sky UI: Click the **Traffic Management** tab and sort the classes by peak bps. Any class that has peaked to more than 50 percent of the link size should be assigned a policy. Make a list.

5. Using your three lists of applications those impacting others' performance, those that are important, and those that peak at more than 50 percent of link size characterize your traffic, categorizing on the basis of importance, size, and sensitivity to latency and jitter.

6. Define partitions and/or policies to protect your important traffic classes. As resources, you can use:

   - Policy/Partition Guidelines (guidelines for choosing appropriate policies and partitions for specific types of traffic)
   - Suggested Policies (on-line, in-product policy suggestions)
   - Reserve Bandwidth for each of the applications you want to protect.

- Insulate users of an application from others using the same application.

7. Define partitions and/or policies to contain your important traffic that impacts others or bursts to a large portion of the link. As resources, you can use:

   - Policy/Partition Guidelines (guidelines for choosing appropriate policies and partitions for specific types of traffic)
   - Suggested Policies (on-line, in-product policy suggestions)
   - Contain a Greedy Business Application (instructions for limiting an important application)

8. Define partitions and/or policies to contain your unimportant traffic that impacts others or bursts to a large portion of the link. As resources, you can use:

   - Policy/Partition Guidelines (guidelines for choosing appropriate policies and partitions for specific types of traffic)
   - Suggested Policies (on-line, in-product policy suggestions)
   - Limit an Application's Total Bandwidth (generic instructions for limiting with a partition)
   - Block Unwanted Traffic (instructions for completely stopping one type of traffic)

Block Unwanted Traffic

Instructions to prevent traffic of a certain type from passing

Steps:

1. Create a traffic class (in Sky or Advanced UI) to identify and isolate the traffic you want to block, if one does not already exist.

   For example, to block attempts to Telnet into a site, create a class called BadTelnet under the Inbound parent class with Telnet as the service in its matching rule.

   For background information, see Traffic Classification Overview and Traffic Tree Overview (Advanced UI, the Sky UI).

2. For web traffic, decide if you want to redirect the user to an alternate web page. The alternate page could simply inform the user of the block. Note the URL of the alternate web page.

3. Determine if the traffic you want to block is UDP or runs over UDP. (You will use a different type of policy to block UDP traffic.) Consult any industry-standard chart of the TCP/IP protocol suite to help in your determination. As examples, Simple Network Management Protocol (SNMP), RADIUS remote authentication, and H.245 Voice over IP (VoIP) are all examples of traffic that runs over UDP.

4. Set a policy on your traffic class to block the traffic:

   To block UDP traffic or traffic that runs over UDP, set a discard policy (in Sky or Advanced UI) on your class.
   To redirect web traffic, set a never-admit policy (in Sky or Advanced UI) on your class using the **web-redirect** option and specify the alternate URL.
   To block web traffic without redirection, set a never-admit policy on your class using the **web-refuse** option.
   For non-web TCP traffic, set a never-admit policy on your class.

   For example, to block incoming Telnet, set a never-admit policy for the BadTelnet class.

   For background information, see Policy Overview.

Block a URL Category

Instructions to prevent users from accessing all websites from a particular URL category

> **Note:** Because PacketShaper gives higher priority to flow delivery and packet processing than to URL categorization, the never-admit policy cannot be used as a comprehensive content filtering feature; it's possible that some requests for websites in blocked categories may not actually get blocked. For a more robust content filtering solution, you would want to add a ProxySG to your deployment.

Steps:

1. Create a traffic class for the URL category you want to block, if one does not already exist. For example, to block user attempts to view porn sites, create a class called Pornography and identify the Pornography category.

   or

   Create a class that classifies multiple categories of similar content (such as all the categories that contain adult content).

2. Decide if you want to redirect the user to an alternate web page that appears when a category is blocked. For example, this web page might explain the company's web-access policy. If you want to do this, make note of the URL of the alternate web page; you will enter this URL in the next step.

3. Set a never-admit policy on your traffic class to block the traffic:

   To redirect users to a web page, set a never-admit policy on your class using the **web-redirect** option and specify the alternate URL.

   To block traffic without redirection, set a never-admit policy on your class using the **web-refuse** option.

4. (Advanced UI only) Repeat the above steps for the Outbound direction.

   > **Note:** the Sky UI has the option to create classes in both directions, eliminating the need to repeat the steps.

5. Make sure shaping is enabled.

Contain a Greedy Business Application

Procedures to manage bandwidth allocation for applications that are important but not urgent, and tend to swell to consume more than their fair share of bandwidth

Some of the types of traffic that fall in this category include FTP, email with large attachments, and print. Although this traffic needs efficient and predictable performance, it can't be allowed to impact the performance of urgent, interactive, and mission-critical applications.

The following procedures are all appropriate to contain bandwidth-hungry, non-urgent applications, but the first is essential:

- Enforce a bandwidth limit on the total of all the greedy, non-urgent traffic you want to control.
  For example, limit all file transfer traffic to 10 percent of WAN capacity.

- Subdivide capped bandwidth for subsets of the greedy, non-urgent traffic.
  For example, if you already limited email to 15 percent of your capacity, you could then reserve half of that 15 percent for email to or from the Sales department.

- Insulate users of the same application from each other.
  For example, if you already limited print traffic to 10 percent of capacity, you could then prevent one high-capacity user from taking most of it with a single print job, leaving others waiting.

Reserve Bandwidth

Instructions to reserve bandwidth for an application or any traffic class that tends to suffer insufficient performance or lose the battle for bandwidth to more aggressive applications

You can reserve bandwidth and therefore protect performance with a PacketShaper partition. For example, you could reserve 30 percent of your capacity for SAP.

Although a partition does indeed protect an entire application or traffic class, it does not enforce any particular performance standards on each session or each user of an application. A rate policy is the appropriate construct for that goal.

Procedure:

1. Create a traffic class (in Sky or Advanced UI) to contain the traffic you want to protect, if one does not already exist.

   Or, alternatively, create a traffic folder (in Sky or Advanced UI) and move several related traffic classes (in Sky or Advanced UI) that you'd like to protect together into the folder.

   Remember to make sure you have corresponding classes in both Inbound and Outbound branches of the traffic tree.

   For background information, see Traffic Tree Overview (Advanced UI, the Sky UI) and Traffic Classification Overview.

2. Decide the minimum amount of your WAN capacity you want to devote to your traffic class. Remember that if the bandwidth is unused, PacketShaper automatically loans it to other traffic in need, so it's never wasted.

   For example, if you have a T1 or E1 WAN link and want to reserve 30 percent for SAP, you could choose 450 Kbps for each direction.

   Decide the maximum amount of your WAN capacity that you want your traffic class to be able to access.

   Some applications, such as Microsoft Exchange, need both protection and limits.
   For help in determining appropriate minimum and maximum amounts of bandwidth, see Sizing a Static Partition.

3. Create a partition to reserve capacity for the total of all traffic in your traffic class. In the **Size** field, enter the minimum you determined in the previous step. If you want the partition to use available excess bandwidth when needed, select the **Burstable** checkbox and enter the maximum you determined in the previous step.

Limit an Application's Total Bandwidth

Instructions to limit bandwidth allocation for applications that tend to swell to consume more than their fair share of bandwidth

For example, contain file transfers to 15 percent of WAN capacity. Or contain music downloads to 2 percent of capacity. Or contain Microsoft Exchange to 40 percent of capacity. Or contain medical digital images to 25 percent of capacity. Or, for service providers, contain one bandwidth customer to the contracted amount.

Additional Information

- The fact that traffic requires a bandwidth limit does not imply anything about its importance or lack thereof. It says only that it tends to consume bandwidth to the extent that others' performance is impacted.

- This procedure does not ensure that an individual session will stay within limits (use a rate policy for that), only that the aggregate of all sessions stays within limits.

Procedure:

1. Create a traffic class (in Sky or Advanced UI) to identify and isolate the traffic you want to contain, if one does not already exist.

   For example, if you want to cap file transfers, then you probably don't need to do anything because the presence of FTP traffic automatically prompts PacketShaper to create FTP classes. But if you want to limit file transfers from a certain server, then you'd need to create a child FTP class that classifies FTP traffic for this server (for example, a class called Inbound/FTP/ServerName).

   For more information, see Traffic Tree Overview (Advanced UI, the Sky UI) and Traffic Classification Overview.

2. Decide the maximum amount of your WAN capacity you are willing to devote to your traffic class for each direction. Remember that if the bandwidth is unused, PacketShaper automatically loans it to other traffic in need, so it's never wasted.

   Decide the minimum amount of your WAN capacity that your traffic requires, even during times of peak congestion. It is not uncommon to decide on a minimum of zero (0 Kbps) for traffic that is to be contained or limited. Minimums that are greater than zero are typically used to protect traffic.

   Both these sizes will be used to define a partition to contain your traffic. For help in determining appropriate sizes, see Sizing a Static Partition.

3. Create a partition to contain the total of all traffic in your traffic class. Make the partition burstable, and use the minimum and maximum from previous steps for the partition's field values. Remember to create partitions for your traffic on both your Inbound and Outbound traffic trees.

Subdivide Bandwidth

Instructions to divide the bandwidth already defined for an application (or other traffic class) among subordinate subsets

For example, if you already limited file transfers to 15 percent of your capacity, you could then reserve half of that 15 percent for the Sales department's transfers.

Note that this is not the appropriate recommendation for dividing bandwidth into equal-sized portions for different users (see Provision Bandwidth Equitably).

Steps:

1. Make sure you already have a traffic class with an associated partition defined.

2. Create a traffic class (in Sky or Advanced UI) for the subset you want to contain or protect. Define it as a child of the class to which you already assigned a partition.

   For example, create a traffic class to identify the subset of file transfers that are initiated by anyone in the Sales organization's subnet. Define a SalesFTP class under the FTP class.

   For background information see Traffic Tree Overview (Advanced UI, the Sky UI) and Traffic Classification Overview.

3. Determine what percentage of of the parent class' bandwidth you want to reserve for this traffic subset. This minimum serves as protection for the traffic subset.

   For example, you could reserve 30 percent of the FTP bandwidth for the Sales group.

4. If you're going to use an explicit bits-per-second size for the child class' partition (rather than a percentage), view partition information to make sure your minimum from the previous step does not exceed the parent class' partition minimum.

   For example, it wouldn't make sense to give SalesFTP 256 Kbps if FTP, its parent, only had 128 Kbps.

5. Determine the maximum amount of the parent class' bandwidth you want the subset to be able to access.

   For more information about choosing partition sizes, see Sizing a Static Partition.

6. Create a partition for your child class to protect or contain its traffic. In the **Size** field, enter the minimum you determined in step 3. If you want the partition to use available excess bandwidth when needed, select the **Burstable** checkbox and enter the maximum you determined in step 5.

7. Repeat the entire procedure for any additional subset of your parent class that you want to manage separately.

Insulate Users of the Same Application

Instructions to protect users of one application from the performance impact of their peers

Not only do users need performance protection from bandwidth-hungry applications, they also need protection from other users of their same application. Consider the difference between these two scenarios:

- A Telnet user experiences poor performance because the total volume of all FTP file downloads is too high.
- An FTP user waits forever because a high-capacity user is downloading a huge file using FTP.

These instructions solve the second scenario, insulating all users of a single application from each other. For a recommendation for the first scenario, see Protect Critical Application Performance.

Procedure:

1. Determine the maximum amount of bandwidth you'd like *each* session of the application to be able to access. For example, you might decide that each FTP session can access no more than 30 Kbps.
2. Determine an appropriate priority for this application or traffic class.
3. Set a rate policy on the traffic class you want to control. Use 0 guaranteed, with the priority and limit values you determined from the previous steps.

Provision Bandwidth Equitably

Instructions to allocate and protect equal amounts of bandwidth for each user

Sometimes organizations care more about equitable bandwidth allocation than about performance for a particular application. Consider these scenarios:

- An educational institution wants each dormitory student to receive a minimum of 20 Kbps and a maximum of 60 Kbps to use in any way he/she wishes.
- Shared-facility buildings (apartments and office buildings) want all tenants to share bandwidth equally.
- An ISP wants to prevent users from using more than their paid-for share of bandwidth, offer bandwidth equity, and juggle oversubscription situations.

PacketShaper, if requested, can create individual subpartitions on the fly as users initiate traffic of a given class. Subpartitions greatly simplify administrative overhead and allow efficient oversubscription. As always, unused bandwidth is available to others in need.

Steps:

1. Create a traffic class (in Sky or Advanced UI) for the traffic you'd like to divide equitably, if one doesn't already exist.

   If you want to give each user a set amount of bandwidth independent of application, all you need are the Inbound and Outbound classes.

   If you want to give each Oracle user, for example, an equal amount, make sure you have an Oracle class.

   And if you want to allocate equitable bandwidth for users of groups of applications, create folder classes, as needed, and move the classes for the applications of interest under the folder class.

2. Determine the amount of bandwidth you'd like users to share. Perhaps it's the entire link size, or perhaps it's a portion of your link. Make sure your size is sufficient to accommodate the anticipated number of users. For example, perhaps you have 1.5 Mbps to be divided among users. You anticipate about 75 users, giving about 20 Kbps each.

   For more information on choosing appropriate sizes for a partition, see Sizing a Static Partition.

3. Create a partition for the traffic class (or folder class) you created earlier. The size of this partition should be the size you determined in the previous step the bandwidth to be divvied up among users.

4. Create a dynamic partition (in Advanced UI only) within the parent partition you just created. The dynamic partition automatically generates small subpartitions that manage per-user or per-subnet bandwidth under the larger, parent partition.

   Consider your plan, using the questions below and their setting suggestions to help guide your choices when defining your subpartition:

Questions and Scenarios	Suggested Dynamic Subpartition Settings
Do you want to provide equitable bandwidth to each user?	Create a subpartition per: Single address
Or to each group of users on the same subnet?	Create a subpartition per: Subnet
Do you want all users to share all available partition bandwidth equitably with no other restrictions?	Subpartition size: 0 Burstable: checked Limit: (blank)
Do you want each user to have access to a capped amount of bandwidth and no more? For example, a service provider might want to cap each user according to purchased plan. A cap of 50 Kbps for each Bronze user, a 150 Kbps cap for each Gold user, and a 250 Kbps cap for each Platinum user.	Subpartition size: 0 Burstable: checked Limit: 50 Kbps (or whatever the maximum size you want)
Do you want each user to have a predetermined amount of bandwidth, no more and no less? For example, a school might want to provide precisely 40 Kbps to each student. If you do use a minimum size as in this example, set a maximum number of users and (optionally) an overflow partition.	Subpartition size: 40 Burstable: not checked Limit: (blank)
Do you want to cut off the number of users at a predetermined number? For example, you could limit the number of users to 100.	Maximum # users: 100 (or whatever the maximum number of users you want to allow)
Or do you want to just keep adding users, cutting the bandwidth proportionately for each? For example, if a 1 Mbps parent partition gets 200 users instead of the anticipated 50, then each user would be accommodated, but they would all have access to only 5 Kbps.	Maximum # users: (blank)
Do you want to refuse latecomers (those who need bandwidth after your maximum is exceeded)?	Overflow partition: (blank)
Or do you want to accommodate latecomers by squeezing them into an overflow partition?	Overflow partition size: <the amount of the parent partition you want to devote to latecomers when there is a surge in demand>

Complete example:

An organization generously decides to give unsanctioned music downloads (such as KaZaA, Gnutella, et al) 5 percent of their T3 WAN link, or 2.25 Mbps. They don't want one high-capacity user stealing the whole thing and leaving their other musical users without resources. But they also don't want to restrict one high-capacity user if he or she is the only active music fan at the moment.

The administrators create a traffic class using matching rules for all the popular music-download applications. They create a static parent partition of 2.25 Mbps. Then they define a dynamic partition with settings as follows:

- Subpartition per: Single address
- Subpartition size: 0 (to distribute the bandwidth evenly among takers)
- Burstable: checked (allow each subpartition to grow to consume available parent-partition bandwidth)
- Limit: (blank) (don't cap any subpartition arbitrarily)
- Maximum # users: 150 (preserve at least 15 Kbps for all users)
- Overflow partition: (blank) (just refuse anyone after 150 active users)

150 users = 150K, burstable to 2.25M each

Quarantine Bandwidth Abusers

Instructions to detect and quarantine those users who use more than their fair share of bandwidth

There are many approaches to keeping applications and users' bandwidth usage in line. Many are covered elsewhere in PacketGuide. For example, you can limit each user with a dynamic partition, cap an application's usage with a partition, or keep each session within reason with a rate policy.

Sometimes you might want to track individual users who use unreasonably large amounts of bandwidth. Once you identify these excessive users, you can enforce limitations on their bandwidth usage without impacting other users. We call this process "placing users into a penalty box." These instructions help you do just that.

With PacketShaper's user awareness feature, hosts are identified by user name, allowing you to readily identify user traffic. You can use Top Talkers or Top Listeners reports to spot these top bandwidth consumers, and then place controls on their traffic usage if they are using inappropriate applications on the network.

Steps:

1. Make sure the user awareness feature is configured.

2. Turn on Top Talkers/Listeners for Inbound and Outbound.

3. Display the Top Talkers/Listeners report for Inbound or Outbound, and look for users who are using an excessive percentage of the link (for example, greater than 12%).

4. To get more insight into what a top user is doing on the network, use the **Add Class** button to create a class for this top user. If you Enable Traffic Discovery when creating the class, PacketShaper will then auto-discover service-based classes for this user. These classes represent the applications the user is accessing. For example:

   Jsmith
   - Gmail
   - WorldofWarcraft
   - Hulu
   - Facebook
   - Default

5. If the top user is accessing primarily recreational applications, you may want to restrict the user's bandwidth usage. To minimize the user's bandwidth usage, assign a policy or partition to the user-based class. For example, you might want to create a partition that caps bandwidth at 64 Kbps.

Quarantine Bandwidth Abusers - Use Host Accounting

The following scenario is used as an example throughout this procedure's steps:
Catch each user (host) moving over one gigabyte of data in a single day. Restrict the user to a maximum rate of 64 Kbps for two days. Then remove the restrictions.

> **Note:** The host accounting feature defines hosts by IP address, not user name. If you want to control bandwidth access for Active Directory users, see Quarantine Bandwidth Abusers.

1. Create a traffic class under **Inbound** to contain your overly exuberant users. Name it **PenaltyBox**. Set protocol = IP, service = any, and server location = any. Create a host list called **abusers** and reference it in the class' matching rule as an inside host list (assuming your users are located on the inside).

   Once the host list and traffic class are defined, change the class' type from Standard to Exception.

   Repeat the same process to create a **PenaltyBox** class under **Outbound**. But this time, you can simply reference the **abusers** host list you already created. You needn't create a second host list.

   Later, you will use these **PenaltyBox** classes to limit the bandwidth that each quarantined user can access. And you will use the **abusers** host list to add and remove users from the PenaltyBox's restrictions.

   For more information, see Traffic Tree Overview (Advanced UI, the Sky UI), Traffic Classification Overview and/or Matching Rules Overview.

2. Create a static partition for your **PenaltyBox** classes with size = 0 Kbps, burstable, and a limit of the maximum bandwidth you'd like to devote to all the abusers together, perhaps 5 to 15 percent of your link size.

3. Define a dynamic subpartition for your **PenaltyBox** classes giving each user the restricted amount of bandwidth you consider appropriate for their penalty phase. For example, you might cap each user at 64 Kbps.

   For more information, see Partition Overview.

4. Enable host accounting using the command-line interface:
   host accounting enable inside 10 3000000

   PacketShaper's host accounting feature, only available with the command-line interface, offers a way to track byte counts of throughput for each IP address, group of users in a host list, or subnet.

   For more information, see host accounting.

   The CLI commands used to this point are as follows:

```
hl new abusers
 class new inbound PenaltyBox inside list:abusers
 class set inbound/PenaltyBox exception
 partition apply inbound/PenaltyBox 0 15%
 partition dynamic apply inbound/PenaltyBox per-address inside 0k 64k
 class new outbound PenaltyBox inside list:abusers
 class set outbound/PenaltyBox exception
 partition apply outbound/PenaltyBox 0k none
 partition dynamic apply outbound/PenaltyBox per-address inside 0k 64k
 host accounting enable inside 10 3000000
```

5. Choose an interval of time over which you'd like to measure bandwidth usage. Perhaps a day or perhaps a week. Each time your interval elapses, you'll need to quarantine new abusers and free previous abusers. For our example, we'll use a day.

   Choose an amount of data for that interval that you consider unreasonable and above normal usage expectations. For our example, we'll consider it unreasonable for any user to fetch or send more than one gigabyte in one day.

6. Let your time interval pass, giving your organization's users enough time to create a usage track record.

7. Retrieve your per-user usage data. Both CLI commands measure dump and host accounting retrieve work to access these measurements. You might want to restrict the report data to a specific subnet(s) or host list to monitor only those IP addresses you really want to examine. For example, universities typically examine only those subnets that go to the students' residential facilities.

   Example command lines:

```
me dump host all by element 1d sort bytes
 host accounting retrieve dns all 6/9 08:00 to 6/10 08:00 into usage.txt
```

   Measure dump has a handy "sort" option so that you can find all your top users quickly and automatically. Host accounting lacks a "sort" option so you must take extra steps to massage the output file to group the top users.

8. If you use the "host accounting retrieve" command, put the results into a text file just as the above example does using the file **usage.txt**. Download the results file to to a PC, open it with Excel or another spreadsheet utility, and sort the rows based on the contents of the "bytes" column. Alternatively, you could use the file with any utility or script that would pull out all rows containing a bytes column bigger than your limit.

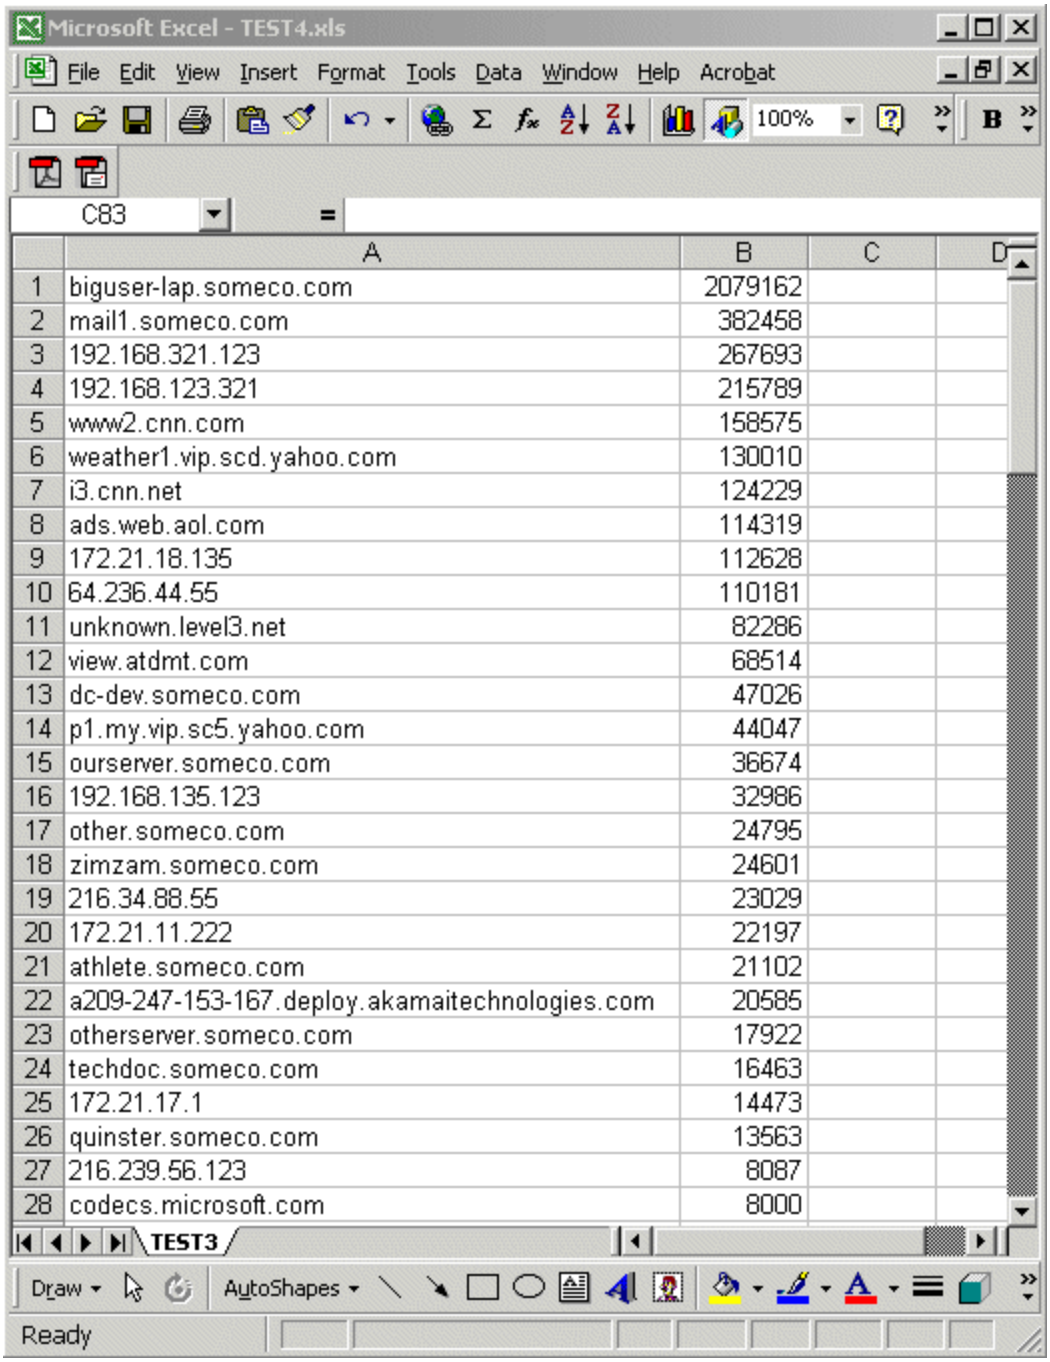
Microsoft Excel - TEST4.xls

File Edit View Insert Format Tools Data Window Help Acrobat

C83 =

	A	B	C	D
1	biguser-lap.someco.com	2079162		
2	mail1.someco.com	382458		
3	192.168.321.123	267693		
4	192.168.123.321	215789		
5	www2.cnn.com	158575		
6	weather1.vip.scd.yahoo.com	130010		
7	i3.cnn.net	124229		
8	ads.web.aol.com	114319		
9	172.21.18.135	112628		
10	64.236.44.55	110181		
11	unknown.level3.net	82286		
12	view.atdmt.com	68514		
13	dc-dev.someco.com	47026		
14	p1.my.vip.sc5.yahoo.com	44047		
15	ourserver.someco.com	36674		
16	192.168.135.123	32986		
17	other.someco.com	24795		
18	zimzam.someco.com	24601		
19	216.34.88.55	23029		
20	172.21.11.222	22197		
21	athlete.someco.com	21102		
22	a209-247-153-167.deploy.akamaitechnologies.com	20585		
23	otherserver.someco.com	17922		
24	techdoc.someco.com	16463		
25	172.21.17.1	14473		
26	quinster.someco.com	13563		
27	216.239.56.123	8087		
28	codecs.microsoft.com	8000		

TEST3

Draw AutoShapes

Ready

9. Examine the top users in your list. Note those whose associated bytes count exceeds your limit.

10. For each user or IP address that you want to quarantine, add the address to your **abusers** host list referenced by your **PenaltyBox** traffic classes. Because you are probably already using the CLI, you might find the CLI command more convenient:
hl add abusers *<the_abuser_host>*

11. To be polite, consider sending the user email to tell them their bandwidth will be restricted for a couple of days due to excessive use.

12. You don't want to leave your quarantined user in the penalty box for too long. Maybe just a brief stay of perhaps two days will prompt the user to reform their habits. After a couple of days, you can edit the host list to remove the quarantined user from the penalty box. It would be convenient if you didn't have to remember to remove the user after two days. The schedule new command allows you to schedule the user's freedom without having to remember later.

    ```
    schedule new once:<mm/dd> <hh:mm> hl delete abusers <the_abuser_host>
    ```

    For example, if today is June 10th, to remove the user *hoggy.yourcompany.com* from the penalty box at 8:00 AM in two days, the CLI command would be:

    ```
    schedule new once:6/12 08:00 hl delete abusers hoggy.yourcompany.com
    ```

> **Note:** You can embellish this procedure in a variety of ways. You can explore using categories in the host accounting feature to restrict the type of usage you examine. You can use the Top Talkers and Top Listeners feature to see who is using the most bandwidth within a given application or class instead of using host accounting (however, you would not get byte counts or be able to save the information for later comparison).

Control Branch Traffic from a Main Site

Instructions to manage the performance of applications to/from multiple branches with a main-site unit

Most PacketGuide content and advice assumes you are managing the traffic moving through your local WAN or Internet link for the benefit of local users. However, frequently a large-capacity unit at a central or main site manages the traffic going to and from multiple branch offices for the benefit of each of those remote branches.

For the most part, the strategy and methods you use to manage only local traffic also apply to managing traffic from multiple branches. You still classify and organize your traffic in a traffic tree; you still analyze utilization and performance; you still protect, contain, pace, and/or provision bandwidth for your applications with all the same commands and windows. The differences center around how you organize your traffic tree and how you allocate an appropriate amount of bandwidth for each location.

When you rely on a main-site unit to manage the traffic and application performance at multiple remote sites, it's best if your WAN is a pure hub-and-spoke design, with all branch traffic going through the main-site PacketShaper. There should be no branch-to-branch traffic, as PacketShaper can manage only the traffic it can see.

Steps:

1. Create a location-based traffic tree that suits your locations' needs and main-site unit's capacities.

   The way your traffic tree organizes traffic dictates how it can analyze and control traffic and the features you have available. For example, if your tree doesn't distinguish traffic to your Paris office from traffic to your Oslo office, then you won't be able to compare traffic volumes or give prescribed amounts of bandwidth to each branch.
   A location-based tree analyzes and controls traffic based on the traffic's source or destination *first*, then, optionally, on other criteria such as application. Choose and create one of the following types of location-based traffic trees:

   - **Simple location-based tree**: manages only the gross bandwidth going to each location. Appropriate for occasions when your primary concern is to provision bandwidth, and you don't particularly care how it's used or how applications perform. In addition, this tree can be handy for a main site when anther PacketShaper sits at each branch location and manages application performance for that location.
   - **Location-based tree with per-location applications**: provides the most power and flexibility in management. Takes more time to configure and usually requires a larger product model due to the number of traffic classes required.
   - **Location-based tree with global applications**: manages each application with the same strategy for all locations. Most difficult tree to understand conceptually, but the easiest to configure and needs the least amount of traffic classes, matching rules, and configuration time. Provides a very scalable and easy-to-manage tree for traffic shaping but does limit visibility.

2. Define a static partition for each branch location's traffic class that corresponds to the size of the branch's WAN link. For example, if you have a branch office in Oslo (with a 512 Kbps WAN link) and a corresponding traffic class called "Oslo," then you should define a static partition of size 512 Kbps.

3. If you have:

- A simple location-based traffic tree
  You're done.

- A location-based tree with per-location applications
  Treat each branch's subtree as an application-based traffic tree, creating appropriate partitions and policies for the application classes as needed generally, only the applications that need to be protected or contained. Look to Policy and Partition Guidelines for help in determining appropriate definitions. If you need assistance with the mechanics of creating the partitions or policies, see Create a Rate Policy (Sky or Advanced UI), Create a Priority Policy (Sky or Advanced UI), Create a Static Partition (Sky or Advanced UI), Create a Dynamic Partition, and instructions for creating other types of PacketGuide policies for assistance.

- A location-based traffic tree with global applications (all branch classes followed by all application classes in your tree), continue with the remainder of this recommendation.

Creating Application Policies for Each Branch's Traffic to Inherit (for location-based traffic trees with global applications)

The goal with this type of traffic tree is to manage any given traffic flow with a partition for the flow's *branch location* and with a policy for the flow's *application*. You already assigned partitions to each branch. Now you'll define policies for the application classes and make sure that each branch's traffic can inherit them.

1. Make each application class inheritable.

   PacketShaper classifies each passing traffic flow in the class that matches the flow's branch location. It adds the metrics for the traffic flow only to the matching branch class. It applies the partition you assigned to the matching branch class.

   But when PacketShaper finds no policy assigned to the branch class, it continues searching until it finds another matching class that is inheritable and has a policy. PacketShaper finds one of your application classes as the second match. The flow inherits the second class' policy. But the second class gets no performance metrics recorded other than an additional policy hit.

   For more on inheritance, see Inheritance Rules.

2. For each application that needs protection and prompt performance, define a rate or priority policy with a priority of 4, 5, 6, or 7 (depending on importance and urgency). Remember that this application will use this policy no matter what branch location it's traveling to or from. For help in determining an appropriate policy for the application, see Policy/Partition Guidelines.

3. For each application that tends to use more than its fair share of bandwidth and undermine the performance of others, define a rate or priority policy with a priority of 1 or 2. For help in determining an appropriate policy for the application, see Policy/Partition Guidelines.

4. For each application that you want to block entirely, define a policy for the traffic class to block the traffic:

   To block UDP traffic or traffic that runs over UDP, set a discard policy (in Sky or Advanced UI) on your class.
   To redirect web traffic, set a never-admit policy (in Sky or Advanced UI) on your class using the **web-redirect** option and specify the alternate URL.
   To block web traffic without redirection, set a never-admit policy on your class using the **web-refuse** option.
   For non-web TCP traffic, set a never-admit policy on your class.

5. After PacketShaper has managed your traffic for some time, click the **Monitor** tab. Notice that, as expected, you have traffic statistics for your branch classes, but not for your application classes. You do have class hits for your application classes. Similarly, you can create utilization reports for your branch classes, but not for your application classes.

Integrate ProxySG with PacketShaper

Instructions on deploying a PacketShaper and a ProxySG at the same site

In a WAN optimization deployment, the PacketShaper and ProxySG appliances can be used at the same site, allowing benefits of both appliances. PacketShaper provides the application traffic visibility and control, while the ProxySG offers a comprehensive WAN optimization solution for distributed locations. See ProxySG Optimization for details on the various ways this appliance can optimize network traffic.

Procedure: Connect the ProxySG

By positioning the ProxySG between the LAN switch and the PacketShaper, the PacketShaper will be able to classify traffic from the ProxySG, provide accurate WAN usage data, and control the optimized traffic going across the WAN.

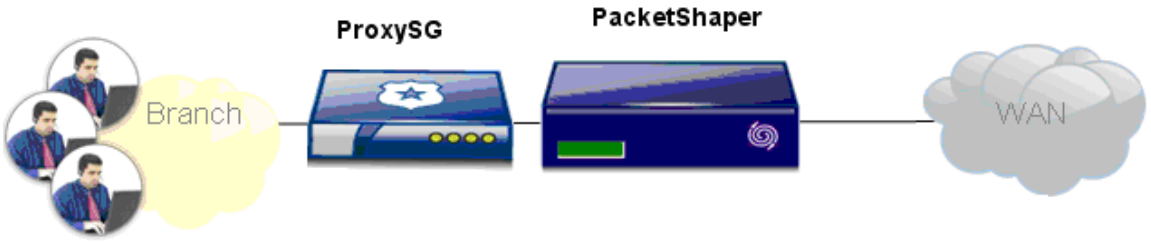

To integrate a ProxySG into a PacketShaper deployment:

1. Disconnect the PacketShaper from the LAN switch.
2. Connect the ProxySG's LAN interface to the LAN switch.
3. Connect the ProxySG's WAN interface to the PacketShaper's INSIDE port.

Procedure: Get Visibility into ADN Traffic

PacketShaper can identify Application Delivery Network (ADN) tunnels and show the type of traffic going through a tunnel.

1. Assuming traffic discovery is enabled and the ProxySG is properly configured for ADN, the PacketShaper will discover ADN tunnel traffic and create classes called ProxySG-ADN in the Inbound and Outbound direction.

2. Go to the Monitor tab to verify the existence of the Inbound/ProxySG-ADN and Outbound/ProxySG-ADN classes.

3. To sub-classify traffic in the ADN tunnel, turn on class discovery for the Inbound/ProxySG-ADN and Outbound/ProxySG-ADN classes.

   When class discovery is enabled, the ProxySG-ADN class will have child classes for the various types of traffic discovered in the ADN tunnel, for example, ProxySG-ADN-FTP and ProxySG-ADN-HTTP. Note that many classes may be auto-discovered, so you may not want to do this if your PacketShaper is nearing its class maximum.

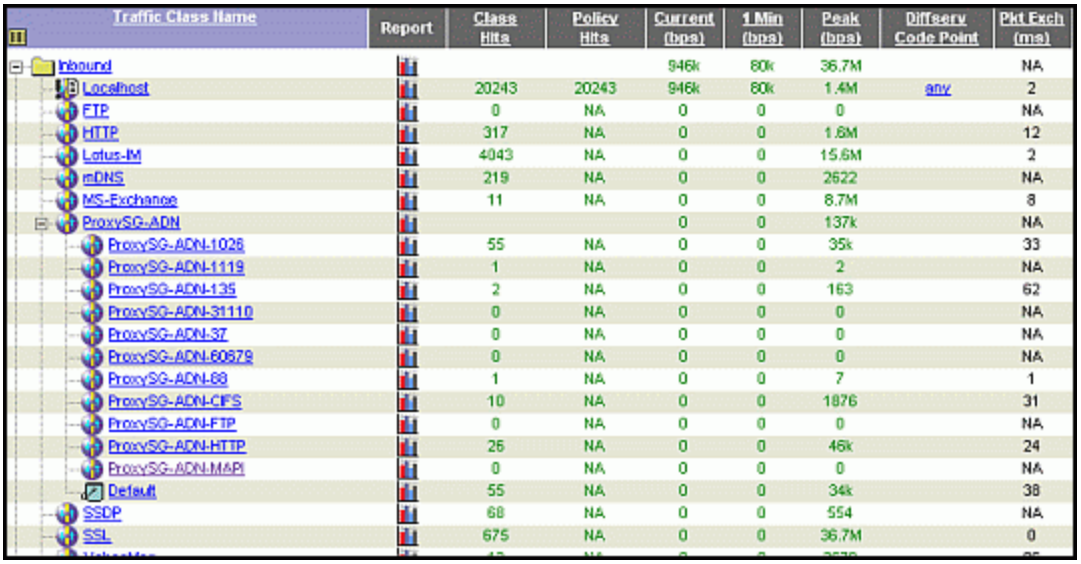

Traffic Class Name	Report	Class Hits	Policy Hits	Current (bps)	1 Min (bps)	Peak (bps)	Diffserv Code Point	Pkt Exch (ms)
Inbound				946k	80k	36.7M		NA
Localhost		20243	20243	946k	80k	1.4M	any	2
FTP		0	NA	0	0	0		NA
HTTP		317	NA	0	0	1.6M		12
Lotus-IM		4043	NA	0	0	15.6M		2
mDNS		219	NA	0	0	2622		NA
MS-Exchange		11	NA	0	0	8.7M		8
ProxySG-ADN				0	0	137k		NA
ProxySG-ADN-1026		55	NA	0	0	35k		33
ProxySG-ADN-1119		1	NA	0	0	2		NA
ProxySG-ADN-135		2	NA	0	0	163		62
ProxySG-ADN-31110		0	NA	0	0	0		NA
ProxySG-ADN-37		0	NA	0	0	0		NA
ProxySG-ADN-60679		0	NA	0	0	0		NA
ProxySG-ADN-88		1	NA	0	0	7		1
ProxySG-ADN-CIFS		10	NA	0	0	1876		31
ProxySG-ADN-FTP		0	NA	0	0	0		NA
ProxySG-ADN-HTTP		26	NA	0	0	46k		24
ProxySG-ADN-MAPI		0	NA	0	0	0		NA
Default		55	NA	0	0	34k		38
SSDP		68	NA	0	0	554		NA
SSL		675	NA	0	0	36.7M		0

Procedure: Disable Flow Limit Policies for ADN

If reflect client IP is not enabled on the ProxySG, there will be many flows from a single IP address (the ProxySG's). The high number of flows from a single client could cause PacketShaper's policy flow limit feature to kick in: PacketShaper will drop any new connections exceeding 10,000 flows for a specific source IP address. In this case, it is recommended to disable policy flow limits for the ProxySG-ADN classes.

To disable policy flow limits for the ProxySG-ADN classes:

> **Note:** If reflect client IP is enabled, you can skip these steps.

1. Access the PacketShaper CLI.

2. Use the policy flowlimit command to disable flow limit policies for the ADN classes. Type the following commands at the CLI prompt:

```
policy flowlimit inbound/proxysg-adn none
 policy flowlimit outbound/proxysg-adn none
```

Procedure: Control Traffic

Using PacketShaper's traffic shaping features, you can control traffic that goes through the ADN tunnel as well as other traffic that goes through the PacketShaper.

To control traffic:

1. *(Optional)* Create a partition (in Sky or Advanced UI) for the tunneled traffic and allocate a maximum percent of overall bandwidth for optimized traffic.

2. To control the traffic the ProxySG cannot optimize, assign policies: rate (Sky or Advanced UI), priority (Sky or Advanced UI), discard (Sky or Advanced UI), never-admit (Sky or Advanced UI). For example, assign a rate policy to VoIP.

3. Enable shaping.

Recommendations for Per-Application Management Strategies

Procedures to manage the performance of specific applications, such as Citrix and peer-to-peer

Manage Citrix Performance

Instructions to differentiate Citrix-based applications and protect their performance

Dependability, centralized management, and cost savings are some of the reasons people invest in Citrix thin-client/server-based architecture. However, unmanaged network traffic frequently destroys the performance of even the most sophisticated server-based computing environment. Citrix-based applications battle for bandwidth with print traffic, web browsing, large email attachments, and other bandwidth-hungry applications.

Tutorial: Managing Citrix Performance (requires Flash player)

Steps:

1. Enable traffic discovery and let PacketShaper discover your Citrix-based applications and create a traffic class for each.

   For example, if you are running PeopleSoft, Microsoft Exchange, and Internet Explorer over Citrix, your traffic tree might look something like this:

   - Citrix
   - Citrix-PeopleSoft
   - Citrix-Exchange
   - Citrix-IExplorer
   - Default

   For background information, see Traffic Tree Overview (Advanced UI, the Sky UI) and Traffic Tree Overview.

2. Create additional traffic classes (in Sky or Advanced UI) for Citrix applications that have not yet appeared on the network, if you need them.

3. If any of your Citrix applications generate a copious amount of print traffic that undermines interactive performance, then divide a single Citrix published application into two traffic classes one for print and one for the rest. You can use Citrix tagging as the means to identify an application's print traffic (a tag value of 3 means print).

   For example, to isolate PeopleSoft's print traffic, you'd expand your traffic tree to look like this:

   - Citrix
   - Citrix-PeopleSoft
   - PSPrint
   - Default
   - Citrix-Exchange
   - Citrix-IExplorer
   - Default

   To accomplish this division, select the Citrix application to be the new parent traffic class (Citrix-PeopleSoft, in the example). Then create a new child class, naming it something appropriate with the application name and the word "print." Choose **Priority 3** from the matching rule's Citrix-specific classification criteria to identify print. The default class for non-print traffic is created automatically. You can change the name to be more appealing, if you like.

4. Create a partition for the Citrix traffic class to carve off a portion of bandwidth to be available any time Citrix applications need it. The partition size should be at least sufficient to cover the number of users you expect with minimally acceptable performance. Consult Sizing a Static Partition for help in determining the partition's correct minimum and maximum sizes.

   Of course, you'll need to balance the performance needs of other, non-Citrix, mission-critical applications and divide resources appropriately. Remember that bandwidth reserved with partitions is never wasted. If unused, it's loaned to other applications in need.

5. Set rate policies for each child class under Citrix to preserve equitable per-session performance. Use a bandwidth rate that ensures the minimum acceptable per-session performance for the Guaranteed field. Cap the session with a reasonable limit to protect other sessions. Assign a priority that matches the urgency of the application.

   For example, the Citrix traffic tree referenced earlier on a T1 link might have the following partition and policy assignments:

   - Citrix (partition: size 400 Kbps, burstable to entire link)
   - Citrix-PeopleSoft
   - PSPrint (rate policy: 0 guaranteed, 20 Kbps limit, priority 2)
   - PSDefault (rate policy: 10 Kbps guaranteed, 30 Kbps limit, priority 6)
   - Citrix-Exchange (rate policy: 20 Kbps guaranteed, 40 Kbps limit, priority 4)
   - Citrix-IExplorer (rate policy: 5 Kbps guaranteed, 30 Kbps limit, priority 3)
   - Default (rate policy: 0 Kbps guaranteed, 15 Kbps limit, priority 3)

   You can stop here with a perfectly fine Citrix solution. But if you want additional control, continue with the next steps.

6. If needed, you can create additional partitions to subdivide the bandwidth for Citrix among its child classes, the published applications. Usually, this isn't necessary, and rate policies for each child class suffice. One circumstance that might require these hierarchical partitions is when two equally critical, urgent, Citrix-based applications are used concurrently and tend to impact each other.

   For example, suppose you have Oracle running over Citrix in addition to PeopleSoft, and you want to give them each the same rate policy. You could create partitions for both Oracle and PeopleSoft giving each, for example, one third the rate you gave to Citrix. (The rest of the Citrix classes would share the other third without requiring further action from you.)

7. It's frequently a good idea to contain applications that tend to consume more than their share of bandwidth in addition to performing the previous steps to protect Citrix traffic. Use the following recommendations to help implement this containment goal:

   - Identify Performance Saboteurs spot those troublesome applications that are undermining others' performance.
   - Limit an Application's Total Bandwidth limit the particularly bandwidth-hungry traffic classes to a specific bandwidth cap.

Balance DICOM's Bandwidth Demands

Instructions for how to manage a deluge of digital-image traffic

Modern MRIs, X-rays, ultrasounds, and CT and PET scans are stored and moved as digital images. Studies can be huge. Typical MRIs contain 10 files of 20 to 30 megabytes (MB) each. Digital images swell to consume most of the bandwidth in limited-capacity WAN links without regard for other applications. Multiple locations, separate imaging facilities, and remote clinicians exacerbate the problem, as images must traverse WAN links multiple times.

DICOM (Digital Imaging and Communications in Medicine) is the global industry standard for transfer of radiological images such as MRI, CT, PET, ultrasound, and mammography. DICOM enables heterogeneous diagnostic and therapeutic equipment and systems to exchange images.

DICOM images are important and must traverse the healthcare network dependably and consistently. However, the images can't be allowed to dominate the network and impose poor performance on all other traffic, including equally important but much more time-sensitive applications. Effectively managing DICOM is a balancing act.

Steps:

1. Identify and differentiate digital image traffic.

   Enable traffic discovery. PacketShaper automatically identifies DICOM, NetBios-IP, NFS, FTP, and RSH, all commonly used in image transfers.

   If your image transfers don't use DICOM or another autodiscovered type of traffic, create your own traffic class (in Sky or Advanced UI) using your own criteria (for example, port number or server). Make sure you have traffic classes for your image traffic in both the Inbound and Outbound branches of the traffic tree.

   If you want to distinguish one DICOM application from another (if you wanted, for example, to treat their traffic differently), then you can divide your DICOM traffic class into multiple, separate classes for different applications.

   The remainder of this recommendation references DICOM only. If the bulk of your image data is in another class, or if you are managing one component of DICOM, the same concepts apply.

   For background information, see Traffic Tree Overview (Advanced UI, the Sky UI) and Traffic Classification Overview.

2. Analyze the impact of DICOM traffic on your network.

   Look at the Top Ten classes (in Sky or Advanced UI). See if DICOM is there. What percentage of the traffic is DICOM?

   Consult the class tree statistics (in Sky or Advanced UI) to look at DICOM traffic data and get an idea of bandwidth trends. Observe the measurements for current rate, one-minute average, and peak rate. Compare the peaks to other traffic classes.

   Create utilization graphs (in Sky or Advanced UI) for your DICOM classes that cover the last week. Is there sustained or frequent high bandwidth usage?

3. Decide an appropriate amount of bandwidth to reserve for DICOM images even during times of much network demand. This will be the minimum size for your DICOM partition. If DICOM doesn't need this whole amount at any time, the excess will be loaned to other traffic (it's never wasted). And if DICOM is very active while there is little demand for other applications, DICOM will be able to access more bandwidth than this minimum indicates.

   The minimum you choose depends on the nature of your other traffic, the volume of your other traffic, your image volume, and other factors. An appropriate size might be 20 percent of your network capacity.

   For background information, see Sizing a Static Partition.

4. Determine a maximum size for your DICOM partition, probably the whole link. DICOM will be able to access this additional bandwidth depending on the presence and urgency of other traffic. After all, you probably don't want to contain DICOM to a minority percentage of the link if there is little or no other traffic.

5. Create a partition to set the total amount of bandwidth for each DICOM traffic class. Use the minimum percentage you determined earlier, select burstable, and either leave the limit blank or enter the maximum you determined in the previous step.

   For example, you might enter 20%, burstable, blank.
   Remember to create partitions for both the Inbound and Outbound branches.

6. Set a rate policy on each of your DICOM classes. A rate policy will indicate the relative importance of your image traffic so that PacketShaper knows how to distribute excess bandwidth. In addition, your DICOM traffic will gain the benefits of TCP Rate Control and have fewer retransmissions that waste bandwidth.

   In the rate policy, use a guaranteed rate of 0 and a priority of 2 or 3 with no limit.

Protect ERP Performance

Instructions to protect the performance of an ERP application

ERP applications frequently lose the battle for bandwidth with file downloads, print traffic, web browsing, large email attachments, and other bandwidth-hungry applications. The result? Slow and unpredictable performance for these business-critical applications.

Follow these instructions to protect the performance of SAP, JD Edwards, or any other ERP application. If your ERP application is Citrix based, see Manage Citrix Performance instead.

Steps:

1. Enable traffic discovery and let PacketShaper discover your ERP application and create a traffic class. If it doesn't appear on your network in a timely fashion, or you want to use your own matching-rule criteria for the traffic class definition, then create your own traffic class (in Sky or Advanced UI).

   For more ideas on classifying specific traffic, see Classification Hints and Examples.
   For background information, see Traffic Tree Overview (Advanced UI, the Sky UI) and Traffic Tree Overview.

2. Create a partition for your ERP traffic class to carve off a portion of bandwidth to be available any time the ERP application needs it. The partition size should be at least sufficient to cover the number of users you expect with minimally acceptable performance. Consult Sizing a Static Partition for help in determining the partition's correct minimum and maximum sizes.

   Of course, you'll need to balance the performance needs of other mission-critical applications and divide resources appropriately. Remember that bandwidth reserved with partitions is never wasted. If unused, it's loaned to other applications in need.

3. Set a rate policy for your ERP traffic class with 0 guaranteed, burstable, no limit, at a high priority, relative to other applications (6, for example).

4. It's a good idea to contain applications that tend to consume more than their share of bandwidth in addition to performing the previous steps to protect traffic. Use the following recommendations to help implement this containment goal:

   - Identify Performance Saboteurs spot those troublesome applications that are undermining others' performance.
   - Limit an Application's Total Bandwidth limit the bandwidth-greedy traffic classes to a specific bandwidth cap.

Common Performance-Management Strategies for Healthcare Environments

Pointers and suggestions on how to manage the performance of applications that are common in the healthcare industry

Networks in the healthcare industry are under siege: more traffic and more technologies going more places. Healthcare organizations must balance the network demands of large digital radiology files, electronic patient records, web-based telemedicine applications, and a variety of other medical and business applications.

Browse the list of common issues and descriptions below. For all that sound like your environment, follow the management suggestions. If the instructions for a topic are short, then they're embedded right here in this file. But if they're longer or covered elsewhere, a link leads to another PacketGuide Recommendation.

- Digital image transfers monopolize the network.

  Without adequate control, large image files dominate WAN links to the detriment of most other applications. Digital images are not good network citizens, consuming as much bandwidth as possible without regard for urgent, interactive applications. Your goal is to keep images moving at a pace that still leaves bandwidth for others: Balance DICOM's Bandwidth Demands.

- Urgent applications using the HL7 standard to exchange information have frustratingly slow response.

  HL7 (Health Level Seven) is an ANSI-accredited organization whose standards facilitate the exchange of clinical and administrative data between many healthcare applications. For example, as an emergency room awaits the results of a lab test before initiating treatment, the lab finishes the test and relays results via HL7. An ambulance communicates with a hospital over satellite to retrieve critical patient information using HL7.

  HL7 flows are very critical, particularly time-sensitive, and typically small. They tend to get crushed under the bandwidth demands of larger traffic flows that aren't necessarily urgent, resulting in inconsistent and sluggish performance.

  Let PacketShaper discover your HL7 traffic automatically and create traffic classes, appropriately named **HL7**. Then assign a priority policy of a high priority (6 or 7) to your HL7 classes.

- Important server-based Citrix applications suffer spotty performance.

  The healthcare industry frequently deploys clinical and business applications to distributed employees over a Citrix platform to take advantage of associated centralized management and cost savings. But thin-client/server-based applications can suffer unpredictable or poor performance on a contention-filled network. The solution is to protect Citrix applications, contain unsanctioned applications, and balance important but greedy applications. See Manage Citrix Performance for help.

- Urgent, interactive applications that are critical for patient care or business functions suffer sluggish and inconsistent performance.

  Applications from companies such as Cerner, Keane, Epic, SAP, Oracle, and McKesson need predictable, prompt performance in any healthcare organization. For help protecting the performance of these applications, see Protect

Application Performance.

- Collaborative sessions between remote professionals suffer stutter and static.

  Diagnostic videos, collaborative VoIP (Voice over IP) sessions, online continuing education, and other streaming traffic arrive in sporadic, irritating bursts as their flows speed up and slow down. To give streaming traffic smooth performance and to limit the number of streaming users to those you can accommodate, see Control Streaming Media, or, for voice/video over IP, see Manage VoIP Sessions.

- Recreational traffic undermines other applications' performance.

  Even in the healthcare industry, there are a few users who don't realize that the network and WAN are literally lifesaving tools. Staff can innocently request music downloads or listen to Internet radio and inadvertently monopolize the network for minutes at a time.

  A recreational traffic problem becomes visible because PacketShaper's automatic classification and analysis features not only reveal the unsanctioned traffic but also the associated users. You can either block the traffic or reserve a small percentage of the network for such purposes.

  To see if recreational traffic is undermining your network's performance, see Identify Performance Saboteurs. For help in blocking or containing recreational traffic, see Block Unwanted Traffic and Control Peer-to-Peer Downloads.

Control Instant Messaging

Instructions to control use of instant messaging, especially for environments where it can be disruptive such as the classroom

Passing notes in a classroom has become an anachronism. Today, children forward electronic messages to each other with a variety of instant messaging protocols. This activity can be just as disruptive, if not more disruptive, than old-fashioned notes.

The following instructions detail how to severely restrict the traffic when disruptive, and let it pass when okay.

Steps:

1. To contain all your instant messaging traffic, create a class (in Sky or Advanced UI) based on the Messaging service group. Name the class InstMsg and create the class in both directions (Inbound and Outbound).

   Note: If your class tree already contained individual instant messaging service classes, you will need to delete them (in Sky or Advanced UI).

2. Decide on your management strategy. For the remainder of this example, we'll suppose you want to severely restrict instant messaging from 8:00 AM until 3:00 PM. Otherwise, you'd like it to have access to the network with a low-to-medium priority.

3. Determine the appropriate policies and partitions for your instant messaging traffic.

   If you block your instant messaging entirely, it has a nasty habit of port hopping jumping from port to port, trying to find an avenue that works. Although PacketShaper's Layer-7 classification can usually still spot the meandering messaging as it switches ports, the network overhead involved in initiating and tearing down all those short-lived connections imposes a high price. ***Therefore, the best approach is not to block instant messaging, but to let it proceed at just a trickle.***

   For our example, during classroom hours, a non-burstable partition with a 2 Kbps size would be appropriate.

   For between 3:00 PM and 8:00 AM, a burstable partition with a size and limit of 8 percent and 30 percent, respectively, of the network WAN link would be appropriate. A priority policy at priority 2 would rank the traffic's access to bandwidth above the 8 percent.

4. If you don't need to vary your management strategy according to a schedule, and you just want one consistent strategy in place at all times, then create your partitions and assign your policies to your InstMsg service group classes.

   Otherwise, if you do need to vary your management strategy according to a schedule, continue to the next step.

5. Determine the CLI (command-line interface) syntax to enforce your commands. For our example, the CLI commands would be:

   For classroom hours:
   partition apply inbound/InstMsg 2k fixed

   For off hours:

```
partition apply inbound/InstMsg 8% 30%
policy apply priority inbound/InstMsg 2
```

Another set of the same CLI commands is needed for the same class names under the outbound branch.

For help with partition and priority policy CLI commands, see partition apply and policy apply priority.

6. Create a command file for each group of CLI commands that needs to be executed together.

   For our example, create two command files, called MorningMessaging and EveningMessaging.

7. Test each command file with the run command and check that the configuration details that you intended to change were indeed changed.

8. Schedule the execution of each of your command files.

Control Peer-to-Peer Downloads

Instructions for controlling the performance impact of P2P file sharing

The whole genre of peer-to-peer (P2P) applications that facilitate file sharing have taken a heavy toll on network performance. Many schools, businesses, and other organizations have no desire to issue oppressive and controversial mandates regarding unsanctioned use of the network. However, they do want to maintain control over their networks, uphold acceptable performance for mission-critical applications, and avoid liability issues from the entertainment industry.

Steps:

The following steps help you to identify, cap or block, analyze, and generally manage this potentially disruptive traffic.

1. To control all of your peer-to-peer traffic, create a class (in Sky or Advanced UI) based on the P2P service group. Name the class P2P and create the class in both directions (Inbound and Outbound).

   > **Note:** If your class tree already contained individual P2P service classes, you will need to delete them (in Sky or Advanced UI).

2. Analyze the impact of P2P traffic on your network and application performance.

   Look at the Top 10 pie charts (in Sky or Advanced UI). See if your P2P class is there. What percentage of the traffic is P2P? Is more than 10 percent of your bandwidth being used by the various types of P2P? If so, some control is probably in order.

3. Decide if you want to block or limit your P2P traffic. If you'd like to entirely block it, use the procedure Block Unwanted Traffic with your P2P traffic classes and stop here. Otherwise, continue with the next steps.

   Many organizations have decided on a strategy where they severely restrict Outbound/P2P traffic but are more lenient with Inbound/P2P traffic. The reasoning behind this strategy is that a large amount of Outbound/P2P usually signifies that someone hosts a file-swapping server and is distributing files to others. But Inbound/P2P traffic is primarily for retrieving files rather than distributing them. To some, legal culpability looms more significantly with distribution.

4. Create partitions to limit the total amount of bandwidth for P2P traffic. Remember to create partitions for both the Inbound and Outbound branches.

   The size of your partitions depend on how restrictive you want to be and the relative importance of other traffic. See Sizing a Static Partition for assistance.

5. Set a rate policy on each of your P2P classes to accomplish several goals: to indicate the relative importance of your P2P traffic so that PacketShaper knows how to distribute excess bandwidth; to insulate P2P users from each other so that one high-capacity user doesn't take the whole P2P partition; to gain the benefits of TCP Rate Control and in particular, to reduce transmissions that waste bandwidth.

   In the rate policy, use a guaranteed rate of 0 and a priority of 0 or 1. Use a per-user maximum of anywhere from 2 Kbps

to 100 Kbps, making sure that it allows a reasonable number of concurrent users within the partitions you created in the previous step.

The procedure described here will catch and control most P2P, file-sharing traffic. It might not catch obscure hybrids or mutations of these protocols. But it is usually sufficient to contain 90 percent of the traffic to maintain a useable network. If this is true for you, stop here. If not, follow the procedure described in Prepare for the Next Napster.

Continue to evaluate the effectiveness of your solution by checking the **Top Ten** tab or the Top Ten pie chart regularly.

Limit Employee Access to Recreational Web Applications

Instructions to permit read-only access to a recreational web application.

Depending on your company's policies, you may want to limit employees' access to recreational web applications so that people aren't wasting company time or bandwidth on unsanctioned activities. For example, it might be acceptable for employees to look at their Netflix account on company time, but not play streaming videos. Or you can allow employees to look at their Facebook pages, but not post anything. These controls are possible by creating web application and operation classes, and then applying appropriate polcies.

Steps:

1. Make sure WebPulse is enabled.

2. To see what operations are available for the web application you want to control, display the list of web applications and expand the web application to see which operations are possible; the operations vary by application.

   > **Note:** Not all web applications have operation support.

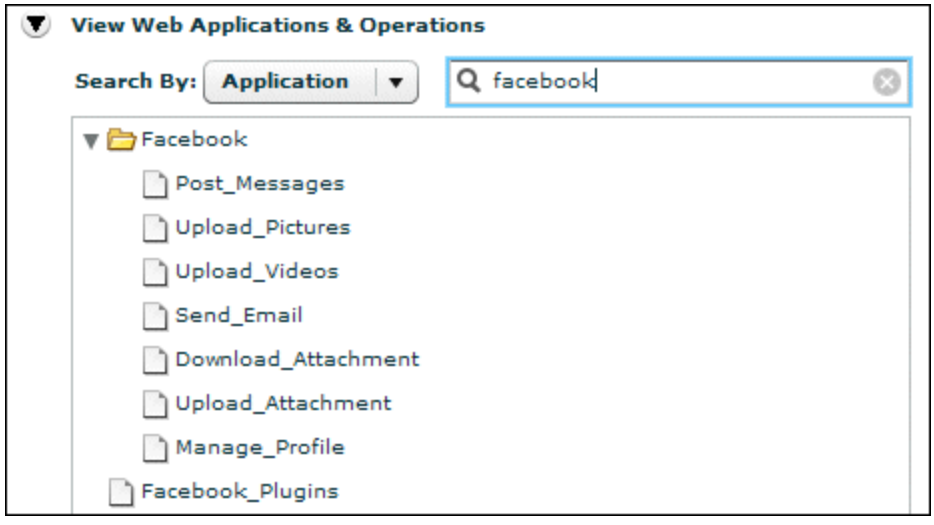

3. If the web application class doesn't already exist, create the class (in Sky or Advanced UI) on both the Inbound and Outbound branches of the traffic tree.

4. Create child classes for each web operation you want to restrict (in Inbound and Outbound). The class should have the

Service set to the web application, and the **Operation** set to the operation you want to restrict. For example, your traffic tree might look similar to this:

Facebook
  Post_Messages
  Upload_Pictures
  Upload_Videos
  Send_Email
  Upload_Attachment
  Default

> **Tip:** If you perform this task in Sky and choose **Both** for the Direction, the classes will automatically be created for Inbound and Outbound. In the Advanced UI, you can create the classes in Inbound and then copy them to Outbound.

5. Apply a never-admit policy (Sky, Advanced UI) to each child class (except for Default).

> **Tip:** If you perform this task in Sky, you can select all the child classes and apply the policy in one operation.

6. Make sure shaping is enabled.

With these policies in place, employees are allowed read-only access to the web application, but cannot do any uploading or posting operations.

Control Users' Access to an Application

Instructions to allow some users access to an application but deny other users' access.

It may be appropriate for some employees in your company to have full access to an application, but inappropriate for other employees to be using the same application on company time. LinkedIn is a good example: recruiters need to be able to send email and post messages to potential candidates, but other employees don't need to have this level of access on the company network.

Here is the Inbound traffic tree for the example used in the steps below:

Inbound
  LinkedIn_Restricted_Operations <-- operations to be restricted
   FullAccess <-- user list
  Default <-- never-admit policy

Steps:

1. Make sure WebPulse is enabled and user awareness is configured.

2. To see what operations are available for the web application you want to control, display the list of web applications and expand the web application to see which operations are possible; the operations vary by application.

   > **Note:** Not all web applications have operation support.

3. In Inbound and Outbound, create a web application class (Sky, Advanced UI) for the application you want to control access to. This class should have the Service set to the web app (for example, LinkedIn) and the Operation set to the first operation you want to restrict (for example, Post_Messages).

4. To this same class, add matching rules (Sky, Advanced UI) for each additional operation you want to restrict.

5. Create a child class of the web app class; this user list class (Sky, Advanced UI) identifies the users who are allowed full access to the web application.

6. Apply a never-admit policy (Sky, Advanced UI) to the Default child class.

7. Make sure shaping is enabled.

With this traffic tree, the users in the FullAccess user list will be allowed to access the restricted operations; all other users will be denied access. Operations that are not part of the restricted class are not classified or controlled in this segment of the tree; if you don't specifically create classes and policies for these operations, all users can access these unrestricted operations.

> **Note:** Until an IP address is associated with a user name, the user is subject to the default policy (never-admit, in this example). Once the user name is identified as one that is on the FullAccess user list, the user will be able to access the restricted operations.

Control Internet Radio Traffic

Instructions for constraining Internet radio traffic

Many employees who are simultaneously listening to Internet radio throughout the business day consume corporate bandwidth and can disrupt your network.

Steps:

The following steps help you to identify and contain this potentially disruptive traffic.

1. Identify and differentiate Internet radio traffic.

   Enable traffic discovery. PacketShaper automatically identifies SHOUTcast, Radio@Netscape, and Windows Media.

   To see if employees are using these applications on your network, go to the Monitor Traffic page and look for Shoutcast, RadioNetscape, and WinMedia traffic classes. If you don't see these classes, either auto-discovery isn't enabled or the applications aren't being used on your network; in either case, you can create the classes manually (in Sky or Advanced UI). Make sure you have traffic classes for your radio traffic in both the Inbound and Outbound branches of the traffic tree.

2. Create a folder class (in Sky or Advanced UI) for your Internet radio traffic and move the radio traffic classes (in Sky or Advanced UI) into this folder. Be sure to create the Radio folder in both the Inbound and Outbound branches.

3. Apply partitions to the Inbound/Radio and Outbound/Radio folders. If your primary concern is protecting business-related, non-recreational applications, you can specify a minimum size of zero (0 Kbps). The Limit should be the portion of the WAN link you are willing to devote to Internet radio.

   The size of your partitions depends on how restrictive you want to be and the relative importance of other traffic. See Sizing a Static Partition for assistance.

4. Set a rate policy on each radio traffic class to reduce jitter.

In addition, you should determine the strategy you'd like to follow if more users want access to Internet radio than your partitions, policies, and link size can accommodate. Would you like to reject latecomers during busy times? Or would you like to squeeze them in with a tiny amount of bandwidth and just string them along until more becomes available? You can use PacketShaper's *admission control* mechanism to define what happens when there isn't enough bandwidth to satisfy guaranteed rate requests.

Control Streaming Media

Instructions to manage the performance of streaming media.

We'll consider streaming media as any content that is *played* (video or audio) as it arrives at its network destination. Streaming media that arrives with stutter and static is not likely to gain many fans. On the other hand, too many fans can undermine performance for everyone, including users of other types of applications. PacketShaper can help keep the equation balanced smooth performance without taking over the network.

PacketShaper's recommendation for managing streaming media entails identifying the streaming-media traffic, reserving a portion of the network to support the media, protecting individual media users from each other, and protecting other applications from the swelling media demand.

For help managing Voice over IP or Video over IP specifically, see the more tailored recommendation Manage Voice and Video Sessions.

Steps:

1. To control all of your streaming media traffic, create a class (in Sky or Advanced UI) based on the Multimedia service group. Name the class Streaming and create the class in both directions (Inbound and Outbound).

   > **Note:** If your class tree already contained individual multimedia service classes, you will need to delete them (in Sky or Advanced UI).

2. Determine appropriate minimum and maximum amounts of bandwidth for streaming traffic. These sizes will be used in partitions that you'll create in the next step.

   The minimum and maximum sizes of your partitions depend on how restrictive you want to be and the relative importance of streaming media with other traffic. For help determining correct minimum and maximum sizes, consult Sizing a Static Partition.

   It is not uncommon to decide on a minimum size of zero (0 Kbps) if your organization does not need streaming applications and your primary concern is protecting other applications. The maximum size, for these types of situations, should be the maximum portion of the WAN link you are willing to devote to streaming media even when there are more sessions than anticipated.

   For environments where streaming-media applications are essential, the minimum size is the portion of the WAN link required for reasonable performance of the anticipated number of sessions. You could allow the partition to burst to the link size by specifying no maximum, or you could specify a maximum that would accommodate peak streaming usage.

3. Create partitions to limit the total amount of bandwidth for streaming applications. Remember to create partitions for Inbound/Streaming and Outbound/Streaming.

4. Set a rate policy on your Streaming classes to accomplish several goals:

   - to indicate the relative importance of your streaming traffic so that PacketShaper knows how to distribute excess bandwidth

- to insulate streaming users from each other so that one high-capacity user doesn't take the whole Stream partition

- to gain the benefits of TCP Rate Control and in particular, to reduce retransmissions that waste bandwidth

- to reserve the appropriate per-session bandwidth to ensure smooth streaming performance

In the policy definition, if you want to preserve smooth per-session performance, use a guaranteed rate of the minimum Kbps for adequate performance. For one of the examples in the previous step, the guaranteed rate would be 15 Kbps.

If you want to protect low-capacity users from high-capacity sessions, cap each session at a Kbps that is more than sufficient for optimal streaming performance. Make sure the limit allows a reasonable number of concurrent sessions within the partitions you created in the previous step.

Assign a priority that communicates the relative importance between the streaming-media application and other applications. It will be used to determine access to the partition's maximum size instead of the minimum.

For example, a rate policy for VoIP could be: Guaranteed: 22K, Limit: 30K, Burstable at priority: 5

5. Determine the strategy you'd like to follow if more users want access to streaming media than your partitions, policies, and link size can accommodate. Would you like to reject latecomers during busy times? Would you like to squeeze them in with a tiny amount of bandwidth and just string them along until more becomes available? Would you like to redirect them to a web page that explains the problem?

   Apply admission control settings to your policies to enforce your strategy.

Manage Voice and Video Sessions

Instructions to manage the performance of Voice over IP or Video over IP streams.

Voice over IP and Video over IP can be managed with similar strategies, as both applications consist of long streaming data sessions accompanied by shorter initiation and control flows. Management recommendations for both Voice over IP and Video over IP are described here and are collectively referenced by the term *V/VoIP*.

PacketShaper's recommendation for managing V/VoIP entails identifying the different types of V/VoIP traffic, reserving a portion of the network to support all of it, speeding the control traffic along quickly, giving each stream an appropriate amount of bandwidth, protecting individual V/VoIP users from each other, and protecting other applications if V/VoIP demand swells.

When instructions pertain to only voice or video, these terms are used. But when the instructions pertain to both traffic types, the all-encompassing term V/VoIP is used. Do not, however, manage voice and video streams together in the same classes with the same policies. Their bandwidth needs are quite different.

Tutorial: Manage Voice over IP Traffic (requires Flash player)

Steps:

1. Create a class (in Sky or Advanced UI) based on the VoIP service group. Name the class VOIP and create the class in both directions (Inbound and Outbound).

   > **Note:** If your class tree already contained individual VoIP service classes, you will need to delete them (in Sky or Advanced UI).

2. To automatically create classes for the VoIP traffic on your network, enable class discovery for the VOIP class (in Sky or Advanced UI) and turn on global traffic discovery.

3. Start and stop a few voice calls or video-conferences to generate some traffic. Leave one of the sessions running at least 10 minutes to collect a substantial amount of measurement data. Note the time interval that your session was running.

   PacketShaper spots many types of V/VoIP traffic automatically, including Clarent, CU-SeeMe, Dialpad, H.323, I-Phone, Media Gateway Control Protocol, H.248/Megaco, Micom, Net2Phone, T.120, SIP, Skype, Vonage, RTP, and RTCP. MCK Communications, Skinny Client Control Protocol (Skinny), and VDOPhone are not auto-discovered but you can manually create classes for them. For a complete list, see Applications, Protocols, and Services.

4. If you have both voice and video, create a pair of folders, one called **Voice** and one called **Video** (in Sky or Advanced UI). These folders allow you to create partitions appropriate for each type of traffic. Be sure to create the folders in Inbound/VOIP and Outbound/VOIP.

5. Move each V/VoIP traffic class into its proper folder under Inbound/VOIP. Do the same for Outbound/VOIP.

   Voice and video clients typically use UDP streams. They have different flows and protocols for initiation, control, and data flows. For example, H.323 starts a conversation on one port (H.323), jumps to another port (Q.931), and eventually splits up into a data flow (RTP) and control flow (RTCP).

6. RTP-I can be further classified by device or codec. For details, see Classify RTP Traffic or Create a VoIP Class in the Sky UI.

   > **Note:** the Sky UI includes a VoIP optimization feature that recommends an appropriate rate policy and partition size for each codec. See VoIP Optimization.

7. Look at V/VoIP traffic data and get an idea of bandwidth trends (in Sky or Advanced UI). Observe the measurements for current rate, one-minute average, and peak rate. Create utilization graphs (in Sky or Advanced UI) covering the time period you left the one session running for the large VoIP data flows (RTP, MiCOM VIP, MCK Voice, or others).

   PacketShaper measures both data payload and header overhead, so requirements will be slightly higher than the nominal bandwidth rate of the protocol. In general, you should expect to see an increase ranging from 10 to 20 percent. Thus, a 768 Kbps stream might take 820 Kbps when headers are included.

8. Determine an appropriate minimum amount of bandwidth you want for your V/VoIP traffic, even during times of contention. In addition, determine the maximum amount of bandwidth you want all V/VoIP to be able to access, even during times of little or no contention. These numbers will be the minimum and maximum sizes for your V/VoIP partition.

   How many concurrent V/VoIP users do you want to support at a minimum? For a rough estimate of bandwidth needs, multiply that number by the amount of bandwidth you observed being used during your one prolonged test session.

   For voice, for example, you might want to support 10 concurrent sessions at 25 Kbps each, for a total of about 250 Kbps in each direction.

   For video, if you're using a 384 codec, for example, you might want to support two concurrent sessions of 420 Kbps each (remember those headers and control information are included).

9. The minimum and maximum sizes of your partitions depend on how restrictive you want to be and the relative importance of V/VoIP with other traffic. For general help determining correct minimum and maximum partition sizes, consult Sizing a Static Partition.

10. Create a partition for your V/VoIP class under both Inbound and Outbound using your numbers from the previous step.

11. Set a priority policy with a relatively high priority (5, 6, or 7) on each of your traffic classes for V/VoIP's setup or control flows (RTCP, SIP, Megaco, MGCP, Skinny, MCK-Signaling, et al).

12. Set a rate policy on each of your V/VoIP data classes (RTP or others) to accomplish several goals:

    - to indicate the relative importance of your streaming traffic so that PacketShaper knows how to distribute excess bandwidth

    - to insulate streaming users from each other

- to gain the benefits of rate control and in particular, to reduce retransmissions that waste bandwidth

- to reserve the appropriate per-session bandwidth to ensure smooth streaming performance

Use a guaranteed rate of the minimum bits per second that are required for acceptable session quality. Make your policy burstable with a relatively high priority.

Typically, if a manufacturer claims that its voice flow requires 8 Kbps, it will actually need 17 to 21 Kbps due to additional overhead and forward error correction. In addition, it is best to overstate the bandwidth needs of UDP traffic by 15 to 20 percent. If in doubt, try 25 Kbps as your guaranteed rate for voice flows.

As examples, rate policies for RTP:

- Voice: Guaranteed: 25K, Burstable at priority 5
- Video: Guaranteed: 420K, Burstable at Priority 5

13. Determine the strategy you'd like to follow if more users want access to V/VoIP than your partitions, policies, and link size can accommodate. Would you like to reject latecomers during busy times? Would you like to squeeze them in with a tiny amount of bandwidth and just string them along until more becomes available? Would you like to redirect them to a web page that explains the problem?

   Apply admission control settings to your policies to enforce your strategy.

Restrict Skype Traffic

Instructions to discourage Skype traffic on your network.

Skype works differently than other Voice over IP (VoIP) clients: it operates on a peer-to-peer model, rather than the more traditional server-client model. This design allows Skype to route calls through other Skype peers on the network to ease the traversal of Symmetric NATs and firewalls. The result is that an individual's computer and network bandwidth may be used to route the calls of other users. Because of liability and security concerns – not to mention unsanctioned bandwidth usage – many companies opt to block Skype traffic on their networks.

Controlling Skype traffic is a bit tricky, though. Because Skype is an aggressive and adaptive application, if you attempt to deny all Skype traffic, the application senses the block and attempts to work around it. Symantec recommends that you allow Skype to set up a call but restrict the VoIP stream: users may deem call quality unacceptable and hang up, and in some cases, Skype will terminate the low-quality call on its own.

Steps:

1. If the Skype class doesn't already exist, create the class (in Sky or Advanced UI) on both the Inbound and Outbound branches of the traffic tree.

2. Create sub-classes for SkypeCommand and SkypeData (in Inbound and Outbound):
   Skype
   SkypeCommand
   SkypeData
   Default

3. Create a partition for the SkypeData class under both Inbound and Outbound. The partition should be a 10 Kbps non-bursting partition.

4. Set a rate policy on each SkypeData class. The rate should be set at a small value such as 1 Kbps.

5. Set a never-admit policy on each SkypeCommand class.

> **Note:** If the Skype client has been configured with a peer/server port that matches one of the port numbers used for PacketShaper's port-based classification, the flow will get classified into a port-based class instead of Skype. In this situation, Skype clients will be able to connect with Skype servers even if you have created a never-admit policy on the Skype class.

Prepare for the Next Napster

Instructions to prevent a new, bandwidth-greedy application from engulfing the network and undermining critical performance before managers have a chance to develop a strategy for dealing with it

When Napster first hit university networks, administrators were caught by surprise. Although many solved Napster's performance impact and legal issues, they were left wondering "What's next?" Most network managers have no desire to restrict an application just because it's new and unknown. But sometimes, on some networks, in some situations, that's just what's most appropriate.

The following procedure describes techniques to contain (not block) new applications. It implements a strategy of "guilty until proven innocent," when it comes to bandwidth allocation. All new, unknown applications are thrown into one bucket and capped. If a new application is deemed sanctioned and appropriate, it can be moved out on its own to share with other sanctioned traffic on a first come, first served basis, using the default policy; or, you can set an appropriate bandwidth-allocation policy.

These same techniques also work for situations when you want to explicitly manage a few applications with custom bandwidth-allocation strategies, and leave the remainder with a default, best-effort strategy.

The following steps cover both types of situations, calling out divergent instructions where needed.

Steps:

1. Enable traffic discovery for PacketShaper as a whole.

2. Tailor your traffic classification tree so that you have explicit classes for all known and supported types of traffic. Most classes will be created automatically by the discovery process. But you'll undoubtedly want to customize the classes a bit.

   If your goal is to explicitly manage only a few applications, then make sure you have the traffic classes for those.

   For background information, see Traffic Tree Overview (Advanced UI, the Sky UI) and Traffic Classification Overview.

3. Create a traffic class (in Sky or Advanced UI) called TraceEffort (or Unsanctioned, Suspicious, Else, OtherIP, or whatever you like) under both Inbound and Outbound in your traffic tree. Its matching rule should contain criteria to match all IP traffic (use **IP** as your **service** and **protocol**). It will serve as a parent class to all IP traffic that does not match your explicit classes. It sits beneath your explicit classes, just above the Default class for Inbound (or Outbound).

4. Move any traffic classes (in Sky or Advanced UI) into TraceEffort that you don't want to explicitly manage or that you consider unsanctioned (depending on your purpose for following this procedure).

5. Determine the minimum and maximum amount of bandwidth you'd like to devote to all traffic under TraceEffort. In addition, determine the priority for its access to the maximum. (You will use these values in step 6 when you create a partition.) Your figures will vary not only with capacity, but with your motivation for choosing this management strategies. Consider these examples with a T1 or E1 link to help you decide:

- A business wants to give all games, music downloads, and Internet radio no reserved bandwidth, a maximum of 10 percent of capacity, but only if absolutely nothing else needs that bandwidth. Min = 0 Kbps; Max = 150 Kbps if using explicit sizes or 10% if using relative sizes; Priority for max = 0.

- A college wants to prevent any performance-impacting surprises, catch any new applications, and prevent them from undermining network performance. But it doesn't want to starve an application just because it's new. They decide to reserve 10 percent of the link for this traffic, cap it at 35 percent, and give it a middle priority of 3.

For background information, see Sizing a Static Partition.

6. Create a partition for the TraceEffort traffic class. Make the partition burstable, and use your minimum and maximum for the partition's field values.

7. Set a rate policy on TraceEffort's Default traffic class. It's the last class indented under TraceEffort. Choose 0 Kbps guaranteed and your priority from a previous step.

   All of TraceEffort's traffic classes will inherit Default's rate policy if they don't have one of their own. By setting this policy, you're giving the classes the efficiency benefits of TCP Rate Control as well as determining the urgency of their access to excess bandwidth.

   For background information, see Inheritance Rules.

Security Recommendations

Procedures to address security issues such as DoS attacks and viruses

Security Best Practices for PacketShaper S-Series

Recommendations from Symantec on best practices for running the PacketShaper appliance securely

PacketShaper offers a variety of methods to make the appliance and its user interfaces secure. The methods range in complexity, security level, and requirements for third-party products. Of course, if your LAN is protected by a firewall, then the unit is protected to the same extent that the rest of your hosts are protected. In addition, you can choose any of the additional methods described in this guide .

Secure Interfaces

Symantec recommends that the PacketShaper appliance be deployed in a secured network that appropriately restricts access to the appliance management network interfaces. This will prevent non-authorized users from accessing the appliance management interfaces—the Symantec Sky UI, Advanced UI, command-line interface, and SNMP. Network administrators can use firewall rules to restrict access to the PacketShaper appliance IP address.

Users with touch access to PacketShaper can secure an interface by creating a list of authorized IP addresses and subnets. Any IP address not on the list or in the subnet range will be blocked access to PacketShaper. For example, you can limit the number of people who can log in to the appliance by specifying just their IP addresses on an inside list.

When the inside, outside, and MGMT interfaces are all set to **secure**, access to the unit is available only via a console connection; the browser interface is disabled.

For instructions on how to secure PacketShaper interfaces, see " Specify Security Settings " on page 297.

Invisible IP Address

PacketShaper appliances are transparent in the network and do not appear in routing tables or as a traceroute hop, so their IP addresses are hidden. A malicious attempt at intrusion would find it much more difficult to guess or discover a PacketShaper unit than a typical router. You need not take any action to gain this type of security.

Virtual, Non-Routable IP Address

If the combination of your firewall and passwords are not sufficient security, you can assign your PacketShaper unit a non-routable virtual IP address. This is a made-up address that is not part of your normal, routable IP address ranges. A PacketShaper unit can recognize its own IP address as traffic passes, even if that address is not a member of the physical subnet where it resides.

Local User Accounts

PacketShaper allows two user roles: the *look* user has read-only access and the *touch* user has read-write access. Touch users can access all commands and full PacketShaper functionality. Look users cannot access commands that modify the PacketShaper's configuration. Nor can the look user retrieve sensitive information, issue commands that would impact the performance of the unit, or create, edit, or delete classes, policies, or partitions. When a look user is logged in, the browser

interface only presents the options that are available with look access. For example, the **class** menu is not shown on the Manage tab because look users cannot add, delete, or rename classes. The look user can view all the setup pages, but the **apply changes** button is not available because this type of user cannot change the settings.

> **Tip:** Only provide a user with the level of access required for his/her role. If someone only needs to view reports and will not be responsible for making configuration changes, provide the user with look credentials only.

To change the passwords, see " Specify Security Settings " on page 297.

> **Note:** Each time you display the Security setup page, the password fields will be populated with eight asterisks, regardless of how long the password is. These asterisks provide extra security in that they prevent anyone from discerning the length of the password string and whether a look or touch password has been set. You do not need to erase the asterisks before applying changes, unless you want to remove the password.

Password guidelines and recommendations:

- Look and touch user passwords can be up to nineteen characters long and are case-sensitive. They can consist of a combination of letters, numbers, and all special characters on the U.S. keyboard.
- The default password for the touch user is *touch* and the look user is *look*. Symantec strongly recommends that these passwords be changed to longer, more complex passwords after initial configuration.
- Be sure to specify different passwords for look and touch users.
- Passwords must be set and cannot be left blank.
- Be aware that passwords are set to their defaults when the unit is reset to factory defaults. Don't forget to redefine the passwords after resetting the unit.

> **Caution:** If you forget your touch password, see "Password Recovery" on page 301.

Third-Party User Authentication

PacketShaper local authentication offers two user accounts: touch and look . Third-party user authentication, such as RADIUS and TACACS+, allow you to have many user accounts and assign different logins/passwords to each user. If you want to take advantage of your third-party user authentication, login, and accounting features, you can integrate PacketShaper with RADIUS or TACACS+.

Secure Unit Access Protocols

By default, only secure protocols are enabled on PacketShaper S-Series. Symantec recommends that you use the industry-standard HTTPS to secure your management sessions with PacketShaper on the browser interface and SSH (Secure Socket Shell) protocol for the command-line interface (CLI).

The following unsecure protocols are *not* allowed on PacketShaper S-Series: FTP, Telnet, and TCP Echo. HTTP and SNMP can be enabled, if necessary.

Supported Ciphers and Algorithms for SSH

PacketShaper supports SSHv2 protocol; SSHv1 is less secure and not supported. However, not all ciphers and MAC (Message Authentication Code) algorithms are supported for SSHv2. A client can select any one of PacketShaper's supported ciphers that are also supported by the client.

Supported Ciphers for SSHv2

aes128-ctr

aes192-ctr

aes256-ctr

aes128-cbc

aes192-cbc

aes256-cbc

Supported MAC Algorithms for SSHv2

hmac-sha2-256

hmac-sha2-512

Unsupported Ciphers

The following ciphers have been disabled due to weakness:

All SSHv1 ciphers (for example: des, des-dbc, blowfish, idea)

aes128-cbc

aes192-cbc

aes256-cbc

aes128-ctr

aes192-ctr

aes256-ctr

arcfour

arcfour128

arcfour256

blowfish-cbc

cast128-cbc

Unsupported MAC Algorithms

The following MAC algorithms have been disabled due to weakness:

hmac-md5

hmac-md5-96

hmac-ripemd160

hmac-sha1

hmac-sha1-96

Transport Layer Security for HTTPS

PacketShaper 11.5.3 and higher versions include the secure cipher suite for Transport Layer Security (TLS) 1.2. The following ciphers are supported:

TLS_DHE_RSA_WITH_AES_128_CBC_SHA256

TLS_DHE_RSA_WITH_AES_128_GCM_SHA256

TLS_DHE_RSA_WITH_AES_256_CBC_SHA256

TLS_DHE_RSA_WITH_AES_256_GCM_SHA384

TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256

TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256

TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384

TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384

TLS_RSA_WITH_AES_128_CBC_SHA256

TLS_RSA_WITH_AES_128_GCM_SHA256

TLS_RSA_WITH_AES_256_CBC_SHA256

TLS_RSA_WITH_AES_256_GCM_SHA384

Starting in PS 11.5.3, PacketShaper uses a self-signed SSL certificate with SHA-256 hash, that uses RSA encryption of a 2048-bit key.

Instead of using PacketShaper's self-signed certificate, you can generate your own PEM file using OpenSSL library, and install the certificate with the **`setup ssl install <pem file>`** CLI command. The PEM file must be stored in the public/3.0/ssl folder.

> **Tip:** Use the File Browser to upload the PEM file to the SSL folder.

If you feel that your PacketShaper's self-signed SSL certificate was compromised, you can regenerate the digital certificate. This task requires that you console connect to the PacketShaper and log in as the Recovery user. You can then choose option **3. Regenerate SSL certificate** from the Recovery menu. You must reboot the system to use the newly generated SSL certificate.

Secure SNMP Monitoring

PacketShaper supports the more secure SNMP version 3. Even though SNMP versions 1 and 2 are also supported, Symantec recommends keeping them disabled to prevent non-authorized users from monitoring the appliance.

SNMPv3 authenticates packets with MD5 or the Secure Hash Al gorithms (SHA), to ensure that a packet has not been tampered with in transit. This feature addresses the need for strong data origin authentication, unlike earlier versions of SNMP (v1 and v2c). SNMPv3's encryption feature provides additional security with Data Encryption Standard (DES) 56-bit encryption to hide the contents of an SNMP packet from unauthorized users. Encryption with AES-128 is also supported. (Note that other key sizes are not supported.) SNMPv3 also offers protection against replay attacks and message stream modification by checking timestamps and message stream integrity.

SNMPv3 encompasses a user-based security model for network management communications. Rather than using the "read" and "write" community strings of prior SNMP implementations, administrators can create specific accounts for each SNMP user and grant privileges through those user accounts. With SNMPv3, each user must belong to a group with a defined security model and security level, as well as a group access policy which specifies which SNMP objects the users can access, and which notifications the group's users will receive.

Secure Configuration Backups

Starting with v11.5.3, PacketShaper encrypts data with the AES-256 cipher and SHA-256 cryptographic algorithm, which provide enhanced security. The PacketShaper configuration settings (including passwords) are stored in the config.ldi file, and can be securely backed up with the **config save** CLI command.

Another way to back up and restore a configuration is with the **setup capture** CLI command. This command creates a command (CMD) file of CLI commands that can be run to recreate a configuration. Because the CMD file is a text file, the **setup capture** output does not include the commands for defining the PacketShaper look and touch passwords.

Session Inactivity Timeouts

For security purposes, users should explicitly log out of PacketShaper:

- To log out of the browser interface, click the **LOG OUT** link in the banner or the **logout** button on the Advanced UI Info tab.
- To log out of the command-line interface, type **exit**.

Logging out discards session content and displays the login page. If a user doesn't explicitly log out, PacketShaper will automatically time out after one hour of inactivity. When a remote login session times out, PacketShaper sends a "timed out" message and disconnects.

Caution: When the Sky user interface has been loaded, the real-time graphs are constantly polling the PacketShaper so the session does not time out.

Multiple failed login attempts do not lock the user out. An individual can continue to try different user names and passwords after authentication failures.

Zeroization Before RMA

Before returning an already configured or deployed appliance to Symantec, the security best practice is to erase all data on the internal persistent storage devices. Erasing persistent data is done using the Recovery utility.

1. Connect the null-modem cable (included with your appliance) from the appliance's CONSOLE port to the serial port on your PC or workstation.
2. Open a terminal emulation program (such as HyperTerminal).
3. Verify that you have configured the program with the following values to communicate with the appliance's CONSOLE serial port:

   9600 bps, 8 data bits, 1 stop bit, no parity, hardware flow control
4. Power cycle the PacketShaper.
5. Log in as the recovery user:

   User name: **recovery**

   Password: **recovery**

   The recovery menu appears. **show screen...**

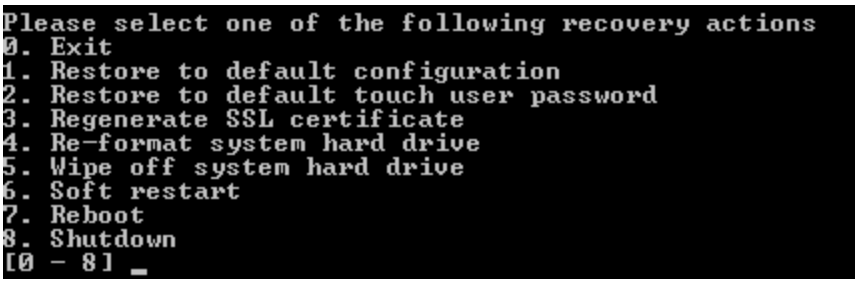

6. Type **4** and press Enter to wipe off the system hard drive.

The hard drive is erased and filled with zeros.

Secure Management by PolicyCenter

When adding PacketShaper units to PolicyCenter, be sure to enable the **Secure Connection** option. When this option is enabled, PacketShaper and the PolicyCenter directory server will use a secure connection (HTTPS) to communicate.

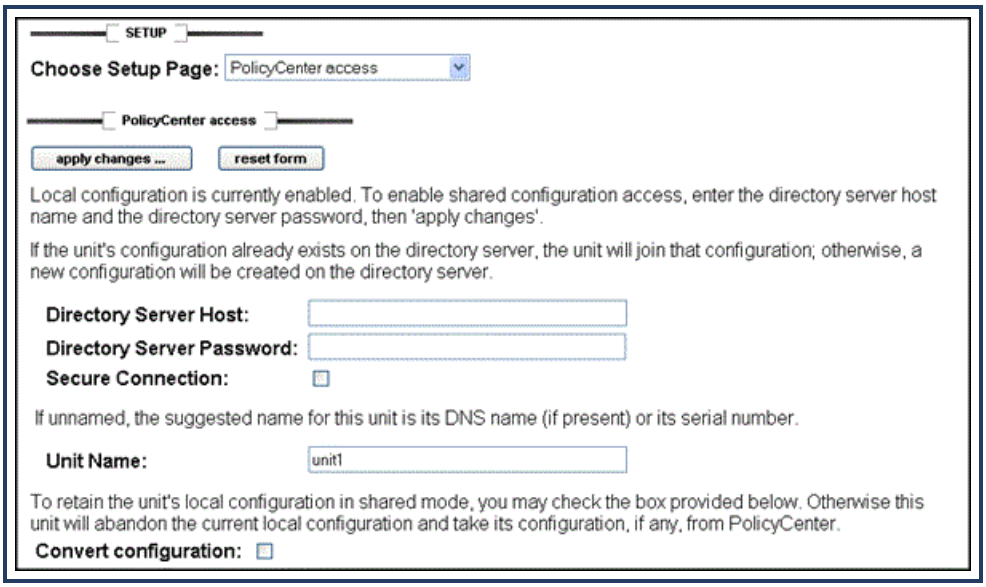
SETUP
Choose Setup Page: PolicyCenter access
PolicyCenter access
apply changes ...
reset form
Local configuration is currently enabled. To enable shared configuration access, enter the directory server host name and the directory server password, then 'apply changes'.
If the unit's configuration already exists on the directory server, the unit will join that configuration; otherwise, a new configuration will be created on the directory server.
Directory Server Host:
Directory Server Password:
Secure Connection:
If unnamed, the suggested name for this unit is its DNS name (if present) or its serial number.
Unit Name:
unit1
To retain the unit's local configuration in shared mode, you may check the box provided below. Otherwise this unit will abandon the current local configuration and take its configuration, if any, from PolicyCenter.
Convert configuration:

Secure Your WAN Applications

Procedures to use the PacketShaper for WAN application security

Traditional IDS/IPS and firewall systems focus on the network perimeter and prevent intrusions from outside hackers and malicious intruders. A firewall uses security policies and filters to prevent unauthorized or potentially dangerous material from entering the network. But what if a threat somehow sneaks through the perimeter and creeps inside the network onto your WAN? Or what if a virus or attack starts inside the LAN?

Companies need security on the inside: every enterprise should have a WAN application security system. This type of system uses the same elements as a perimeter IPS, but applies them to applications and traffic on the WAN or LAN. The PacketShaper offers a number of features that focus on preventing disruptions and unauthorized activities on the network. This recommendation shows you how to use PacketShaper features to perform the following types of network security:

- Application flow validation
- Rogue server detection
- Application overload protection
- Detection of recreational traffic
- Protecting against Network Congestion Events (NCEs)
- WAN application forensics

Application Flow Validation

The growing complexities associated with network traffic make sophisticated classification techniques a necessity. Simple IP address or static port schemes fall short. PacketShaper's classification detects dynamic and migrating port assignments, differentiates applications using the same port, and uses Layer 7 application indicators to identify applications. The PacketShaper allows you to verify that an application flow is what it's supposed to be.

As a response to the PacketShaper's powerful classification and control capabilities, clever hackers continually attempt to circumvent the road blocks PacketShaper has put up for them. With the PacketShaper's ability to identify and block/contain recreational traffic, people have tried different ways to masquerade this traffic as innocent web-browsing.

Web Proxies

Suppose some of your employees have installed web proxy tools (such as proxifier, proxster, and proxyshare) that allow them to send traffic (such as KaZaA, Morpheus, and ICQ) through an HTTP tunnel. Since HTTP traffic (port 80) can get through the firewall, they are able to bypass the firewall rules. However, the PacketShaper is able to spot peer-to-peer traffic hiding inside an HTTP tunnel. The HTTP-Tunnel service automatically identifies and classifies traffic that is sent through an HTTP tunnel via an HTTP proxy server on the Internet. Once this traffic is classified, you can restrict this unsanctioned traffic by applying appropriate policies. (See Control Peer-to-Peer Downloads.)

This technique assumes all HTTP tunnel traffic is unsanctioned. But what if your website is using HTTP tunneling via a web proxy to secure outbound traffic via SSL? How can you limit the use of unsanctioned web proxies but still ensure the sanctioned

ones are protected? The PacketShaper can easily handle both situations. Just create an IP-based class for the sanctioned proxies and set appropriate policies to protect the traffic flowing to and from sanctioned proxies. The unsanctioned web proxy traffic will get classified into the inbound/http-tunnel and outbound/http-tunnel classes that have the restrictive policies.

Encryption

Sometimes peer-to-peer traffic uses encryption to try to disguise itself and avoid detection. When applications, such as Ares, obscure their content through encryption, the traffic can circumvent security. Allowing encrypted P2P traffic exposes organizations to potential worms, viruses, spyware, and other unauthorized files. Because the PacketShaper is able to identify connection setup sessions for P2P applications that utilize encryption to obscure transferred content (such as Ares 1.8.1), you have the ability to block or discard the connections.

Port Hopping

Port hopping is another technique developed for avoiding application detection. Many P2P protocols, such as AOL (America Online) and KaZaA, use port hopping so that they can bypass firewalls and router configurations that are designed to block or rate-limit them. When an application has port hopping capability, it will try to reach a remote host on a different port when it is unable to connect on the default port. For example, a P2P program might first try to connect on port 80. If this fails, it might try on port 3904. If port 3904 doesn't work, it might jump to port 4556. Because the PacketShaper understands protocols up to level 7 of the ISO model, it is able to recognize many port-hopping protocols no matter what port they are using.

PacketShaper automatically validates each flow and places it in the proper traffic class. To see which application (service) PacketShaper classified a flow into, you can use the traffic flow command. To see a complete list of all the applications the PacketShaper is able to classify, see Applications, Protocols, and Services.

Rogue Server Detection

Rogue servers are not authorized, don't contain the latest security patches, and don't have anti-virus software installed. These rogue servers are a network breach waiting to happen, exposing the company to liability, viruses, and attacks.

PacketShaper offers a couple features that can aid in the detection of rogue servers: Top Talkers/Listeners and Flow Detail Records.

Top Talkers and Top Listeners

One way that you may notice that something is amiss on your network is a sudden increase in a certain type of traffic, such as FTP. Suppose FTP is not usually in the top ten applications on your network, and it suddenly appears on the Top Ten pie chart.

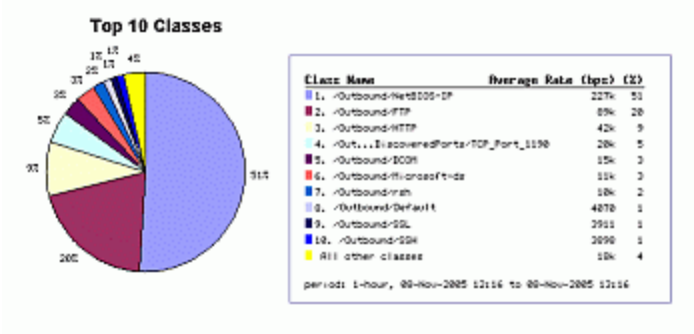

To find out who is contributing to this increase in FTP traffic, you can use the Top Talkers and Top Listeners feature to track the hosts that initiate the most traffic (talkers) and hosts that receive the most traffic (listeners) on the FTP class. You can track top hosts for up to 32 classes at one time.

Flow Detail Records

Flow Detail Records (FDRs) provide detailed traffic information such as flow origin and destination, flow size, throughput, efficiency, and service type. PacketShaper units serve as *emitters*, automatically pushing data to remote systems, called *collectors*. The collectors interpret and incorporate the data into meaningful reports. Symantec's centralized network reporting product, IntelligenceCenter, includes one such collector called DataCollector. DataCollector's FDR capabilities can assist with detailed flow and host analysis.

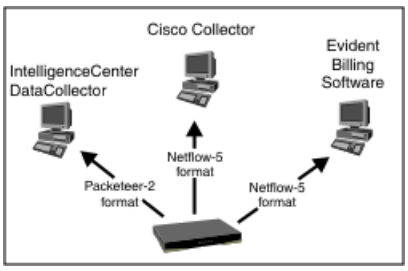

By configuring each of your PacketShapers to emit FDRs to the DataCollector, you have a wealth of reporting and analysis tools available to you. Using FDR, you can drill down to detailed information based on individual host IP addresses. You do not need to enable top talkers/listeners on a class: DataCollector stores these details automatically. This capability helps in the process of detecting a rogue server on your network.

For more information, see Flow Detail Records Overview and IntelligenceCenter Overview.

Application Overload Protection

If a denial-of-service attack occurs on your network, you'll want to make sure your application servers don't become overwhelmed with the excessive traffic. *Flow limit policies* offer such protection. These types of policies limit the number of flows per minute that can be received by any server. You can then use the *adaptive response* feature to notify you when flows are being blocked because of the flow limit policy. While the flow limit policy protects servers from being overwhelmed, adaptive response lets you monitor what's happening and can be set up to automatically send you an email when flows are being blocked by the flow limit policy. To do this, you will want to create an agent based on the Class ME Variables template, the Inbound/Default class, and the server-flood-block class measurement variable.

For more details on using flow limit policies, see the recommendation titled Detect and Limit DoS Attacks. For background information on adaptive response, see Adaptive Response Overview.

Detection of SYN and Spoofing Attacks

A high number of new flows per minute and/or failed flows per minute can indicate an attack is occurring. Adaptive response offers several agent templates that monitor flows to/from hosts to help in the detection of SYN attacks and spoofing. After setting up these adaptive response agents, you can be notified automatically when a host has excessive failed flows and/or new flows per minute, thus alerting you to a possible attack. In addition, you can automatically restrict the bandwidth of these violating hosts so that you can minimize the problem while you do further investigation. For complete details, see Use Agents to Detect SYN Attacks and Use Agents to Detect Spoofing.

Detection of Recreational Traffic

Excessive recreational traffic can disrupt a network. Employees downloading music files, hosting a file server for music files, communicating via unsanctioned instant messaging, listening to Internet radio all day long, watching streaming video of current events – all of these non-business-related activities consume bandwidth that could be better used elsewhere. You can use PacketShaper's sophisticated classification and control features to detect and control recreational traffic.

For procedures, refer to the following recommendations:

Control Peer-to-Peer Downloads

Control Instant Messaging

Control Internet Radio Traffic

Protecting Against Network Congestion Events

The very nature of TCP traffic can cause Network Congestion Events (NCEs) – network-related situations that impact application performance and user productivity. Every day tasks, such as an employee FTPing a large file or emails with gargantuan attachments, can produce NCEs. Fortunately, the PacketShaper can prevent these types of NCEs from occurring.

The basic foundation of the PacketShaper is designed to prevent Network Congestion Events. Partitions can reserve sufficient bandwidth for mission critical applications. Rate or priority policies make sure critical applications get first dibs on the bandwidth they need. Restrictive partitions and policies can limit the bandwidth of unsanctioned traffic. Rate policies allow you to take advantage of TCP Rate Control to smooth out bursty applications so that they don't consume the entire link.

With this infrastructure in place, a large file download will not create network congestion. If half the company views a streaming video of a major news event at the same time, the performance of mission critical applications will not be impacted. Or, if a number of employees simultaneously download the just-released 100MB Microsoft Service Pack, the entire network will not become painfully slow.

WAN Application Forensics

Network forensics is the capture, recording, and analysis of network events in order to discover the source of security attacks or other problem incidents. This term is borrowed from the legal field, where forensics pertains to the investigation of crimes.

The PacketShaper offers several tools to help in network forensics:

- Graphing tools
- Flow Detail Records
- Packet Capture

Graphing Tools

The PacketShaper measures hundreds of characteristics about traffic as it passes, creating an extensive collection of measurement data. PacketShaper stores measurement data on appliances for up to two months. PacketShaper's metrics can also be incorporated into SNMP management platforms, NetFlow v5 collectors and third-party reporting tools, such as Concorde, InfoVista , and Microsoft Excel.

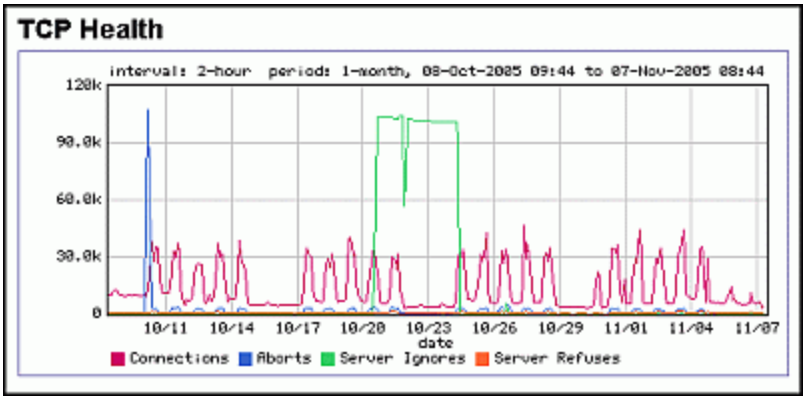

The PacketShaper uses this measurement data to generate a variety of graphs, several of which can assist in the diagnostic process. The TCP Health graph shown above gives you a comprehensive picture of TCP connections for a link, partition, or traffic class. It compares the number of TCP connections that were started, aborted, ignored, or refused by the server. This can help you identify where problems occur. For example, if you see a large number of Server Ignores, you know that you have an overloaded or malfunctioning server that is ignoring new connection requests.

For instructions on creating graphs, see Graph a Specific Link, Partition, or Class.

For descriptions and examples of all PacketShaper graphs, see Preconfigured Graphs.

Flow Detail Records

As mentioned earlier, the Flow Detail Record feature provides drill-down metrics on a per-flow basis that include items such as flow origin and destination, flow size (in packets and bytes), when the flow was sent, the flow's application or service, the flow's Layer-4 protocol and IP ToS/Diffserv bits, the type of controls that were applied to the flow, and more.

This granular level of detail opens up a wealth of opportunity for enhanced troubleshooting and forensic help. For example, you can:

- Examine the "chattiest" host IP pairs for traffic from a specific application, location, or combination of the two.
- Split the traffic from one branch office into its different application, service, or DSCP types, even if you didn't sub-classify traffic into its services as it passed.
- List traffic's busiest ports; list the ports a specific application or host used; list which applications used a specific port; spot potential port scans.
- Expose the top current or historical traffic contributors or recipients for a location or application

If you specify IntelligenceCenter DataCollector as the FDR collector, you can use a variety of reports to aid in troubleshooting network problems and help determine the source of a DoS attack. For example, IntelligenceCenter's flow detail reports let you see the busiest hosts on your network. If you spot a suspicious host, you can drill down to see the applications used and destination addresses contacted.

For more information about FDR, see:

Flow Detail Records Overview

Host Analysis

The PacketShaper maintains a database of all hosts that have active connections through the unit. Once a host closes its connection, the host will be purged from the database. In addition, the PacketShaper will clear host entries if they aren't active for approximately ten minutes. Thus, the host database is a real-time list of hosts.

Host Analysis is a reporting tool that allows you to view, sort, and drill-down on hosts in the host database. If you spot a suspicious host (for instance, one that is using excessive bandwidth or creating an inordinate number of connections or failed connections), you can drill down to find out more detailed information. What other hosts is the suspicious host communicating with? What are the host's current and recent traffic flows? What type of traffic is it (HTTP, SSL, etc.)? The Host Analysis tool can provide the answers. See Perform Host Analysis for details.

Packet Capture

For a more in-depth analysis of your flows, PacketShaper offers the *packet capture* feature. This feature captures packets for future analysis, allowing you to analyze detailed information about the packets in a traffic class, such as the source and destination IP addresses and protocols used. You can even examine the packet down to the byte level, if you are so inclined. With the packet capture tool, you can capture a small sample of traffic from an unknown or troublesome flow so that you can investigate it in detail: view packet headers, real-time display of top users, or content at specific offsets into packets.

First, you need to decide which packets you would like to collect. A major advantage of using the PacketShaper as a collector is that you define precisely which traffic to capture. You don’t have to collect huge log files with mostly irrelevant traffic. Any traffic you can identify with matching rules in a traffic class can be captured independently. For example, if you want to capture all Telnet packets to a certain IP address – you can. Or if you want to capture only Oracle traffic for one particular database – you can.

The basic process is:

1. Specify which traffic class(es) from which to capture packets.
2. Enable packet logging for a period of time.
3. Open the resulting log file in a third-party analysis tool, such as EtherPeek, Ethereal, or a Sniffer.

For more information, see the recommendation titled Sniff without a Sniffer.

Detect and Limit DoS Attacks

A collection of procedures to prevent DoS (denial-of-service) attacks

Recent denial-of-service attacks against popular websites have made headlines and raised consciousness about vulnerability. These insidious attacks employ a variety of mechanisms to wreak havoc. The suggestions below are a collection of techniques using PacketShaper to reduce the impact of a variety of different types of attacks. You can pursue any or all. Keep in mind that PacketShaper products are not firewalls, but they do offer some rather sophisticated help in this particular security arena.

The ultimate targets of DOS attacks are not the only victims. Frequently, if the type of DoS attack involves flooding a target with too much traffic or too many TCP connections, the attack methodology creates countless other victims so-called "launch pads." Launch pads are third-party servers that are forced into being unwitting assailants; they're infected with instructions to generate large amounts of traffic or connections to be sent to the intended final target.

The attack-foiling methods described here:

- Monitor for attacks
- Limit the rate to or from one host
- Limit ICMP floods
- Limit concurrent flows of one type
- Detect pseudo-address attacks
- Disrupt virus attacks

Monitor for Attacks

The techniques in this section show you several different ways to check whether a denial-of-service attack has occurred on your network. The other sections show you more proactive ways to detect and limit DoS attacks.

Monitor class hits.

Class hits let you know the number of times flows match a class. By monitoring class hits regularly, you'll know what constitutes a normal number of hits in a class and an abnormally high number of hits will be apparent.

1. Click the **Monitor** tab. The Monitor Traffic window appears.

2. Look at the numbers in the Class Hits column. These values have been accumulating since the unit was last reset or since the statistics were reset.

3. If desired, reset the monitor statistics and check back later to see the values of the class hits.

4. Repeat step 3 on a periodic basis to monitor class hits.

5. Optionally, click the **Class Hits** column heading to sort by class hits. The classes with the greatest hits will sort to the top of the list.

Spot a huge increase on an ignored traffic class.

An attack can use a traffic type that has been assigned an ignore policy (remember, this policy does not toss traffic but simply ignores it and lets it pass).

1. Select the **Monitor** tab and scan your traffic tree for any classes that are assigned an ignore policy.

2. Check the Class Hits column for classes with ignore policies. Keep tabs frequently enough so that you can spot abnormalities.

   If you spot any suspiciously high class hits, continue to the next steps.

3. Change the policy so that PacketShaper will track statistics for its traffic. (PacketShaper does not keep statistics on classes with ignore policies.) A priority policy with priority 3 might be appropriate.

4. Set Top Talkers and Top Listeners for the class.

5. Turn traffic discovery on within the class.

6. After waiting a few minutes for enough traffic to pass, examine any available information:
   - Look at Top Talkers to see who contributes to the traffic.
   - Look at the utilization graphs. Is it really a flood?
   - Look at the traffic tree. After you turned on traffic discovery for the class, did PacketShaper add any child classes to refine what makes up the traffic? Who are the recipients? Is it legitimate traffic?

Limit Flood Attacks (Back to top)

Flood attacks establish a large number of illegitimate connections, consuming bandwidth or overwhelming hosts. Attacks can originate from single or multiple hosts and be destined for single or multiple hosts.

Limit the Rate of New Traffic Flows to or from One Host

PacketShaper can detect and limit the impact of SYN flood or similar denial-of-service attacks that are directed at a particular host or if the attack is from a specific IP address. In order to improve PacketShaper s ability to automatically contain denial-of-service attacks, PacketShaper can enforce a limit on the number of flows per minute that can be initiated by any client or

received by any server. The limits are set to default values of 10,000 flows per minute on client hosts and 100,000 flows per minute on servers. Flow limits are automatically set on any classes that have a policy set on them, and PacketShaper will block any flows that exceed these limits.

1. Access the CLI.

2. Use the hostdb info CLI command to determine the normal numbers of flows per minute on the client and server sides of a host.

3. If you want to set or adjust the default limits on a particular class, use the CLI command policy flowlimit to restrict the rate at which any given client or any given server can initiate new flows.

   For example, enter **policy flowlimit /Inbound/Default 100000 200000** to restrict the number of new flows from any given client to one hundred thousand per minute and from any given server to two hundred thousand per minute. Depending on network size and conditions, you may find lower values are more appropriate for your network

   Any traffic that is classified in Inbound/Default or is in any traffic class that inherits Inbound/Default's policies is held to these limits. (Background information on inheritance.)

4. If you have any traffic classes that do not inherit the policy you just defined, consider whether you want these classes' clients/servers to be included in the evaluation of your host limits. If so, repeat the **policy flowlimit** command for each class.

5. Remember to apply flow limits to your Outbound traffic as well as your Inbound traffic.

   If one of your hosts has been infected with a virus that initiates copious amounts of traffic to deluge an intended target, you have been turned into a DoS launchpad. You can avoid the legal and service implications of contributing to someone else's attack by using the **policy flowlimit** command on Outbound traffic.

Limit the Amount of ICMP Traffic

Many DoS attacks use the ICMP protocol. Normally, ICMP traffic uses only a tiny portion of bandwidth. There are exceptions, especially if you require large numbers of pings, typical in hosting environments for mail, games, and so on.

1. ICMP traffic is automatically classified by PacketShaper and has its own class. Check the Monitor window for ICMP's peak usage or display ICMP's Utilization graph for an interval over the last month.

   Take the class' highest peak number, double it, and jot it down.

   If you have no historical data to examine (for example, if you're just setting up), you could just take 5 or 10 percent of your link (for example, 75 Kbps). However, this alternative is not advisable if you have an unusually high amount of ICMP on your network in normal conditions. It's best to let PacketShaper gather a week or two of ICMP statistics and use the previous method.

2. Create a non-burstable partition for Inbound/ICMP and Outbound/ICMP traffic classes using the rate you jotted down as the size.

Limit the Number of Concurrent Traffic Flows for One Traffic Class

1. Examine the metric "licenses-peak" over the last month to spot trends for the highest number of simultaneous sessions for a traffic class. Double the highest number and jot it down.

2. Limit the number of allowable concurrent traffic flows for one traffic class using the CLI command class licenses. Use the number you jotted down in the previous step. Note that although rate policies and partitions limit the amount of bandwidth, this command limits the number of simultaneous flows. For example, you could enter **class licenses Inbound/Oracle 5000** on the command line to contain Oracle to 5000 simultaneous sessions.

3. Use the CLI command traffic licenses to monitor the concurrent flows for the class.

Detect Psuedo-Address Attacks (Back to top)

Spoofing and LAND attacks send packets that appear to be from a trusted source by maliciously setting the source and/or destination to false addresses. The traffic can then gain access to hosts or services that should be secure.

Disrupt Virus Attacks (Back to top)

If a virus contains an identifiable string in a field you can specify in a matching rule, then you can use a traffic class definition to spot and toss the traffic. A virus using HTTP with a characteristic sequence of characters in its URL is the easiest traffic to block and a very common vehicle for viruses.

1. Create a traffic class under Inbound/ HTTP to collect all instances of the virus. Use this URL criteria:
   Name: *AnyName*
   Protocol family: IP
   Service Type: HTTP
   Service Location: Any
   Class Criterion: URL : *use known URL with asterisks for wildcards*

   For example, the Code Red virus from months ago used a URL with one of two characteristic strings. Customers used the following line in their class definition to catch the first:
   Class Criterion: URL: */default.ida?NNNNNNNNNNNNNN*

   The Nimda virus from the same vintage used:
   Class Criterion: URL : */root.exe*

2. Add a matching rule using any additional known URL criteria: (*optional*)
   Protocol Family: IP

Service Type: HTTP
Server location: Any
Class Criterion: URL: *use known URL with asterisks for wildcards*

For example, the Code Red virus mentioned in the previous step used the following final line in its additional matching rule:
Class Criterion: URL:*/default.ida?XXXXXXXXXXXXXXX*

And the Nimda virus used:
Class Criterion: URL:*readme.eml*

3. Simulate an attack. From any web browser, enter **www.yahoo.com/*url pattern you used***. You will get a message telling you that the document was not found. Refresh the page. This will show up as a hit on your class (check the **Monitor** tab).

   For example, the nimda virus could be tested using:
   www.yahoo.com/system32/cmd.exe

4. When you are certain the class is collecting only the proper traffic, apply a never-admit policy to the class.

5. Copy the class to the Outbound branch to prevent any existing virus from spreading.

Caution: Although PacketShaper can help prevent this kind of traffic and has some firewall characteristics, it is not a firewall. High traffic demands on PacketShaper can result in improper classification; this is not a risk worth taking. We recommend that you have a firewall in place dedicated to protection. Attacks have been known to get past firewalls under the best circumstances—use PacketShaper as additional protection.

Administration Recommendations

Procedures to address administrative goals

Vary Management Strategy by Time and Day

Instructions to change bandwidth-allocation policies according to the clock or calendar

Occasionally, you might want to control performance differently at different times or on different days. Consider these examples:

- A teacher objects to instant messaging in her classroom during teaching hours, but it's okay during lunch or after school.
- A company's network administrator does not want to permit games or MP3 music downloads on weekdays, but feels that it should be okay on weekends.
- A sales-ordering application needs access to twice its usual bandwidth in the last two days of the month because the sales personnel typically deliver the most orders during those last two days.
- A bank's branch offices must synchronize their databases daily but have very constrained WAN links. They want to allow this synchronization to monopolize the links at night, but not during the day when other traffic must have priority.

You can vary your configuration details based on the day or the time of day. The choice of day can be daily, weekends, weekdays, specific dates, specific days of the week, and/or specific dates of the month.

Steps:

1. Create a traffic class for the traffic you want to control and schedule, if you don't already have one.

   For example, you could create a traffic class called OracleAccounts for Oracle traffic using the Accounts database.

2. Plan your management strategy and the PacketShaper policies or partitions you'll need. You can use Policy/Partition Guidelines for help.

   For example, with a 128 Kbps link, you might want to contain OracleAccounts to 30 Kbps during business hours and allow it to fill the link after hours. And you always want the benefits of TCP Rate Control. You'd need:

   Rate policy with no guaranteed rate, burstable with priority 6. (all the time)
   Static partition of size 20 Kbps, burstable, 30 Kbps limit. (for business hours)
   Static partition of size 30 Kbps, burstable with no limit. (for after hours)

3. For those polices or partitions that you want all the time, simply apply them in your normal way. For example, set a rate policy.

4. For those partitions and polices that you want active only at certain times, determine the CLI (command-line interface) syntax to enforce your commands. For example, consult the CLI command partition apply.

5. Create a command file for each group of CLI commands that needs to be executed together.

   If you have many schedule-based changes, you might consider having command files named WeekdayMornings, WeekdayNights, EndOfMonth, and so on. Then you can insert all the CLI commands for all the changes required at the indicated times.

   For example, suppose you want different policies and partitions for both MP3 downloads and gaming based on weekend versus weekday. You need one command file for MP3s and gaming that contains the weekend polices and partitions, and another command file for the same traffic classes that contains the weekday policies and partitions.

   Remember, you need not include a CLI command in your command file unless it needs changing. All existing policies, partitions, and so on that are in force will remain so unless you include a command imposing change. So, for example, if you include gaming policies in command files for each Monday and Saturday, you needn't include any gaming policies in your command files for morning and night.

6. Test each command file with the run command and check that the configuration details that you intended to change were indeed changed.

7. Schedule the execution of each of your command files.

Migrate Legacy Configurations to PacketShaper S-Series

Instructions to migrate the configuration of a legacy PacketShaper to a PacketShaper S-Series appliance

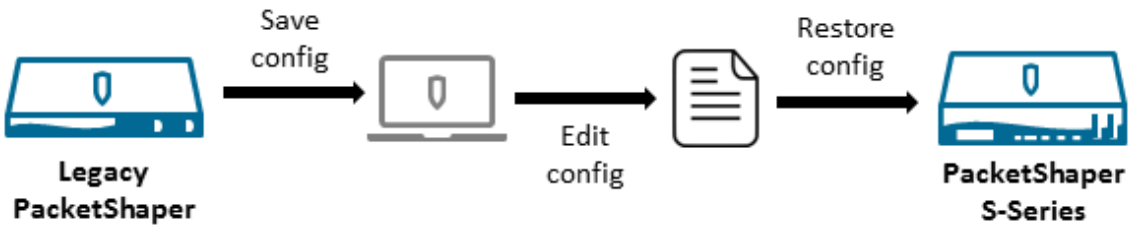

Steps:

Save the Legacy PacketShaper Configuration

First, you must create the configuration files and store them on a local workstation.

1. Log in to the legacy PacketShaper command-line interface.
2. Enter the following CLI commands to capture the configuration:
   `cd cmd`
   `setup capture complete legacy.cmd`
3. (Optional) Enter the following CLI command to save a backup of the configuration:
   `config save backup`
4. Create a zip archive of the legacy configuration files (legacy.cmd, backup.cmd, backup.ldi). See " Create a ZIP File" on page 689.
5. Save the legacy.zip file to a local workstation. See " Download a File" on page 693.

Prepare the Configuration File

Before you put the configuration file on the PacketShaper S-Series appliance, you need to do some prep work on the file to make it compatible with features and components on the S-Series.

1. Unzip the legacy.zip file on your local workstation.
2. Edit legacy.cmd in a text editor, and make the following changes:
   - If the S-Series will use different IP addresses than the legacy appliance, modify the appliance IP address, gateway settings, and/or DNS server.
   - Convert the legacy device names (e.g., main) and port names (e.g., inside) to S-Series names (e.g., slot2, slot2_in). See device and port names for your S-Series model: S200, S400, S500.

- (Optional) Comment-out (precede line with #) any lines for features not available in S-Series: `tunnel`, `wccp`, `frame`, `set ssh`, `set https`, `set ssl`.
- If the configuration has any traffic classes based on device names (e.g., main), you will need to convert them to S-Series device names (e.g., slot2). See device names for your S-Series model: S200, S400, S500.
- Save this file as s-series.cmd. You will be using this file to restore the legacy configuration on the PacketShaper S-Series.

Migrate the Configuration to PacketShaper S-Series

> **Tip:** Because rebuilding the configuration can take some time (especially if the class tree is large), Symantec recommends performing this procedure while the PacketShaper S-Series is offline or during off-peak hours.

1. Use the Quick Start Guide poster to set up the new PacketShaper S-Series.
2. Upload the s-series.cmd file to the new appliance. See " Upload (Copy) a File to PacketShaper " on page 694.
3. Log in to the PacketShaper S-Series command-line interface.
4. If the PacketShaper S-Series had previously been running another configuration, you should clear these settings before restoring the legacy configuration:
   `config reset`
5. Enter the following CLI command to restore the settings from the legacy configuration:
   `run s-series.cmd`

> **Note:** If you didn't comment-out the lines of legacy-only features, the .out file will contain error messages, such as *Error: There is no command 'tunnel'*. These errors can be disregarded as they don't cause issues with the configuration on the S-Series appliance.

Differences in Bandwidth Measurement Between Legacy and S-Series

Legacy PacketShapers measure bandwidth differently than PacketShaper S-Series. In PS 9.2, the Ethernet header is not included in bandwidth calculations, whereas they are included in PS 11.x. As a consequence, partitions that were created when running on legacy PacketShapers may need to be resized when migrated to PacketShaper S-Series. These partitions may need to be increased in size by up to 20% for 64-byte packets; other packet sizes may require different partition size adjustments after migration to S-Series.

PacketShaper S-Series uses the following formula when measuring bandwidth consumed by a frame:

```
(total frame size + linkOverheadBytes) + (total frame size * linkOverheadPpt / 1024)
```

By default linkOverheadBytes and linkOverheadPpt are set to zero so do not affect the frame size. For more information about the linkOverheadBytes and linkOverheadPpt variables, see setup variable.

Verify PacketShaper Health

Instructions to use a PacketShaper to investigate the health and state of the unit itself.

Available Capacity

PacketShaper has several types of capacity limits that can impact your ability to expand your traffic management solution. For example, if you run out of traffic classes or matching rules, you won't be able to auto-discover or add any new types of traffic.

It's a good idea to check your unit(s)' available capacities every few months or whenever you have made lots of configuration changes (such as changes to your traffic class tree, your policies and partitions, the amount of traffic your network supports, and so on).

You can check estimates of system boundaries based on unit model. Keep in mind these are rough estimates of maximum limits that vary with real-time conditions.

You can check your real-time system limits with the CLI command called sys limits.

You'll see results similar to these: **show screen...**

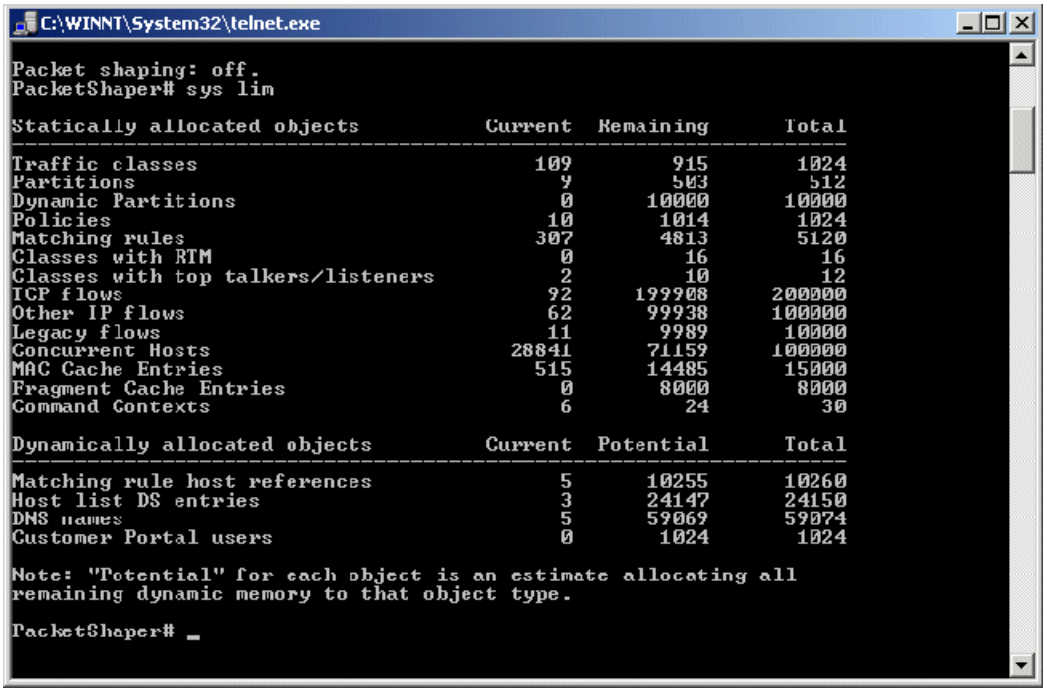

If you find that your are running out of capacity, there are a variety of remedies, depending on which item and which unit model you have. You can:

- Upgrade to a higher capacity model.
- Consider trimming your traffic tree. Often, the traffic class tree gets full because auto-discovery creates a lot of classes for traffic that passed only once. You can go through the class tree and delete these unnecessary classes to free space.
- Contact Symantec's support organization for help.

Diagnostic Commands

Explore some of PacketShaper's other diagnostic CLI commands such as net nic, version, and uptime.

Advanced UI Tasks - Step-by-Step

Step-by-step instructions on how to accomplish a specific task in the Advanced UI

Setup Tasks

When you initially set up your unit and answered the questions asked in Guided Setup, you specified settings for enabling connectivity and access to your network (such as the unit's IP address, subnet mask,and link size). You can modify any of these basic settings, via the Setup tab in the browser interface. In addition, you can configure special features such as WebPulse, security, and user awareness.

Run Guided Setup

After you install the PacketShaper into your network, you need to configure the PacketShaper software.

> **Caution:** DNS must be configured and working correctly for the licensing process to work. Make sure to configure at least one DNS server during Guided Setup.

1. Access Guided Setup via a console connection.
2. Enter the default User Name [**touch**] and Password [**touch**].
3. Answer the questions as they are asked, using the information from the Pre-Installation Checklist.
4. When you are finished, commit your changes. Your configuration will be saved to your appliance.

License PacketShaper

Install Initial License

Upgrade PacketShaper License

Install Initial License

When you purchase a PacketShaper license, the license key will automatically be synced to your appliance within 24 hours, assuming it has Internet access. If your appliance is not connected to the Internet, you will need to download the license key, place it on a web server or workstation that your PacketShaper can access, and then load the license.

> **Caution:** DNS must be configured and working correctly for the licensing process to work. See " Reconfigure Basic Settings" on page 289

Loading the License Key (with Internet Access)

Since the license synchronization may not be immediate, you can initiate the license retrieval process by issuing the **license update** CLI command, as described below.

1. Log in to the PacketShaper CLI. See "Access the Command-Line Interface" on page 1006.
2. Type version verbose to see if the license has been installed.
3. If the message *Invalid license key* appears, issue the license update CLI command.
4. Type reset to reboot the appliance.
5. After the reboot, type version verbose. Verify the license key attributes and expiration dates.

Loading the License (without Internet Access)

If your PacketShaper S-series appliance does not have Internet access, follow the procedure below to manually download the license key from the MySymantec portal, upload the key to your PacketShaper, and then activate your license. For more information on downloading a license, refer to the Getting Started web page.

1. Locate the PacketShaper serial number in your shipping alert email.
2. Log in to your MySymantec account.
3. In the products list, locate the serial number of the appliance you want to license.
4. Generate a new key; this downloads the license file.
5. *Option A:* Put the license key file on a web server that the PacketShaper can access.
    a. Copy the license key file to the web server.
    b. Log in to the PacketShaper CLI. See "Access the Command-Line Interface" on page 1006.
    c. Use the license load CLI command to load the license from the web server. For example:
    license load http://mywebserver.com/license/license.bcl

> **Note:** Filenames are case sensitive. Be sure to enter the license key filename with the same upper/lower case you saved the file with.

d. Skip the next step and proceed to step 7.

6. *Option B*: Put the license key file on the PacketShaper:

   a. Log in to your PacketShaper browser interface.

   b. In the Advanced UI, open the File Browser.

   c. Upload the license key file to a folder on your PacketShaper (such as *9.256/*).

   > **Note:** Make note of the filename as you will need to specify it when loading the license in the next step.

   d. Use the license load CLI command to load the license. For example:
   license load file://localhost/9.256/license.bcl

   > **Note:** Filenames are case sensitive. Be sure to enter the license key filename with the same upper/lower case you saved the file with.

7. Reset the PacketShaper to activate the new license.

   > **Note:** Please be patient as rebooting may take quite some time; the amount of time depends on your network environment. DO NOT unplug the power cables during the reset process as this can cause a corruption of the configuration.

8. After logging back in to the appliance, use the version verbose CLI command to verify the new key attributes and expiration dates.

9. If you loaded a license key that changes the number of supported traffic classes, you will need to reset measurement data. See " Reset Measurement Data" on page 572 or measure reset.

Additional Information

- The license activation and expiration dates appear next to **License Status** on the Info tab in the Advanced and Sky UIs.
- If a license will be expiring within the next 90 days, a message (such as *PacketShaper license will expire in 15 days*) appears in the CLI banner and on the Info tab in the Advanced and Sky UIs.
- If the license expires, the PacketShaper will go into passthrough mode; packet shaping and classification will not function until you load a valid license.

- For software add-ons and upgrades, you will need to enter an activation code to download the license key. See Upgrade PacketShaper License.

Upgrade PacketShaper License

You can purchase license upgrades to extend the capabilities of your PacketShaper S-series appliance. For example, if you have a PS-S500 with a 2G link key, you can upgrade to a 5G or 10G link key.

Create the Upgrade License

1. You can either click the link in your e-fulfillment email, or follow the steps below.
    a. Go to the Symantec Support site, MySymantec.com, and log in using your MySymantec credentials.
    b. In the products list, locate the serial number of the appliance you want to license.
    c. Go to the **Upgrade** tab.
2. Follow the steps below to update the license on the appliance. If your PacketShaper does not have access to the Internet, go to " Upgrade PacketShaper License " above.

Updating the PacketShaper License (with Internet Access)

Since it can take up to 24 hours for the license to synchronize to your PacketShaper, you can initiate the license retrieval process by issuing the **license update** CLI command, as described below.

1. Log in to the PacketShaper CLI. See "Access the Command-Line Interface" on page 1006.
2. Type version verbose to see if the license has been upgraded.
3. If the upgraded license hasn't synched to the appliance, issue the license update CLI command.
4. Type reset to reboot the appliance.
5. After the reboot, type version verbose. Verify the license key attributes and expiration dates.

> **Note:** If the upgraded license still does not appear on the appliance, you can try manually downloading and uploading the license key as described in the next section.

Updating the PacketShaper License (without Internet Access)

If your PacketShaper S-series appliance does not have Internet access, follow the procedure below to manually download the license key from the MySymantec portal, upload the key to your PacketShaper, and then activate your license. For more information on downloading a license, refer to the Getting Started web page.

1. Locate the PacketShaper serial number in your shipping alert email.
2. Log in to your MySymantec account.
3. In the products list, locate the serial number of the appliance you want to license.

4. Generate a new key; this downloads the license file.
5. *Option A:* Put the license key file on a web server that the PacketShaper can access.
    a. Copy the license key file to the web server.
    b. Log in to the PacketShaper CLI. See "Access the Command-Line Interface" on page 1006.
    c. Use the license load CLI command to load the license from the web server. For example:
    license load http://mywebserver.com/license/license.bcl

    > **Note:** Filenames are case sensitive. Be sure to enter the license key filename with the same upper/lower case you saved the file with.

    d. Skip the next step and proceed to step 7.
6. *Option B*: Put the license key file on the PacketShaper:
    a. Log in to your PacketShaper browser interface.
    b. In the Advanced UI, open the File Browser.
    c. Upload the license key file to a folder on your PacketShaper (such as *9.256/*).

    > **Note:** Make note of the filename as you will need to specify it when loading the license in the next step.

    d. Use the license load CLI command to load the license. For example:
    license load file://localhost/9.256/license.bcl

    > **Note:** Filenames are case sensitive. Be sure to enter the license key filename with the same upper/lower case you saved the file with.

7. Reset the PacketShaper to activate the new license.

    > **Note:** Please be patient as rebooting may take quite some time; the amount of time depends on your network environment. DO NOT unplug the power cables during the reset process as this can cause a corruption of the configuration.

8. After logging back in to the appliance, use the version verbose CLI command to verify the new key attributes and expiration dates.
9. If you loaded a license key that changes the number of supported traffic classes, you will need to reset measurement data. See " Reset Measurement Data" on page 572 or measure reset.

> **Note:** If the license expires, the PacketShaper will go into passthrough mode; packet shaping and classification will not function until you load a valid license.

Modify Settings

When you initially set up your unit and answered the questions asked in Guided Setup, you specified settings for enabling connectivity and access to your network (such as the unit's IP address, subnet mask, and link size). You can modify any of these initial settings, via the Setup tab in the browser interface.

To modify settings:

1. Click the **Setup** tab. The basic settings are automatically displayed in the Setup window. **show screen...**

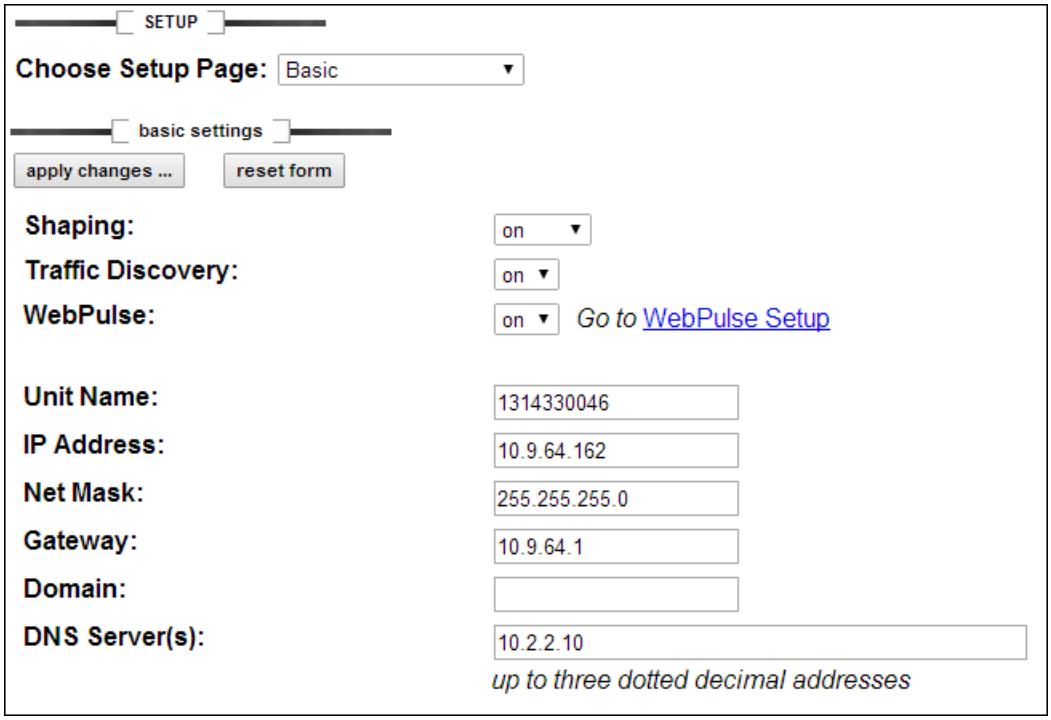

2. Use the **Choose Setup Page** drop-down list to access other setup pages. The following table describes the functions of the available setup pages.

Setup Page	Function
Basic	Set the basic configuration required for deployment on a specific network. See Reconfigure Basic Settings.
Adaptive Response	PacketShaper's adaptive response feature monitors unit, application, and network health, providing a color-coded summary that lets you identify potential problems. See Adaptive Response Overview.
Action File Distribution	View or change adaptive response action file distribution settings. See PacketGuide for PolicyCenter for more information on distributing action files.
Command Scheduling	Schedule commands to be executed at a specific time, on a specific day, on an ongoing basis or only once. See Create a Schedule.
Date & Time	Set date and time; configure time zone and SNTP. See Modify Date and Time Settings.
Email	Configure an SMTP (Simple Mail Transfer Protocol) server for the schedule, Adaptive Response, and user event features. See Configure an SMTP Server.

Setup Page	Function
Flow Detail Records	Configure PacketShaper to emit flow detail records. See Define Flow Detail Record Collectors.
Heartbeat	Configure PacketShaper to emit heartbeats. See Configure Heartbeat Emission.
Image	View or change image prescription settings. See PacketGuide for PolicyCenter for more information on distributing images.
PolicyCenter Access	Configure PacketShaper for use with Symantec PolicyCenter, an appliance designed to manage and centrally control a network of PacketShapers. See "Subscribe PacketShaper Units to PolicyCenter" on page 294. **Note:** PolicyCenter is a separate appliance and must be purchased separately.
RADIUS Client	Enable RADIUS authentication. See " Configure RADIUS" on page 302.
Security	Set access passwords, select allowable access methods, and define a login message. See " Specify Security Settings " on page 297.
Service Groups	Create or modify service groups. A service group is a set of services that have been grouped together based on a common functionality; they offer a quick way to classify and control a group of related services. See "Service Group Tasks" on page 433.
SNMP	Configure Simple Network Management Protocol (SNMP). See " Configure SNMP on PacketShaper" on page 334.
System Variables	Adjust PacketShaper's default settings. See Adjust System Variables.
TACACS+ Client	Enable TACACS authentication. See " Configure TACACS+ " on page 311.
Unit Resets	Reset configurations to factory defaults. See Reset the Traffic Tree, Reset Measurement Data, Reset Unit to Factory Default Settings, Reboot the Unit.
User Awareness	Configure the BCAAA server. See Configure PacketShaper as a BCAAA Client.
Watch Mode	View and configure watch mode settings. See Configure Watch Mode.
WebPulse	Configure the WebPulse cloud service so that PacketShaper can classify URL categories, web applications, and operations. See Configure WebPulse.

Reconfigure Basic Settings

When you deploy a unit on your network, you must configure settings, such as an IP address, to enable communication with the unit. These settings are configured during the Guided Setup process. Periodically, you may need to modify these settings.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To view or modify the basic settings:

1. Click the **Setup** tab. The Basic Settings automatically appear on the *Setup* screen. **show screen...**

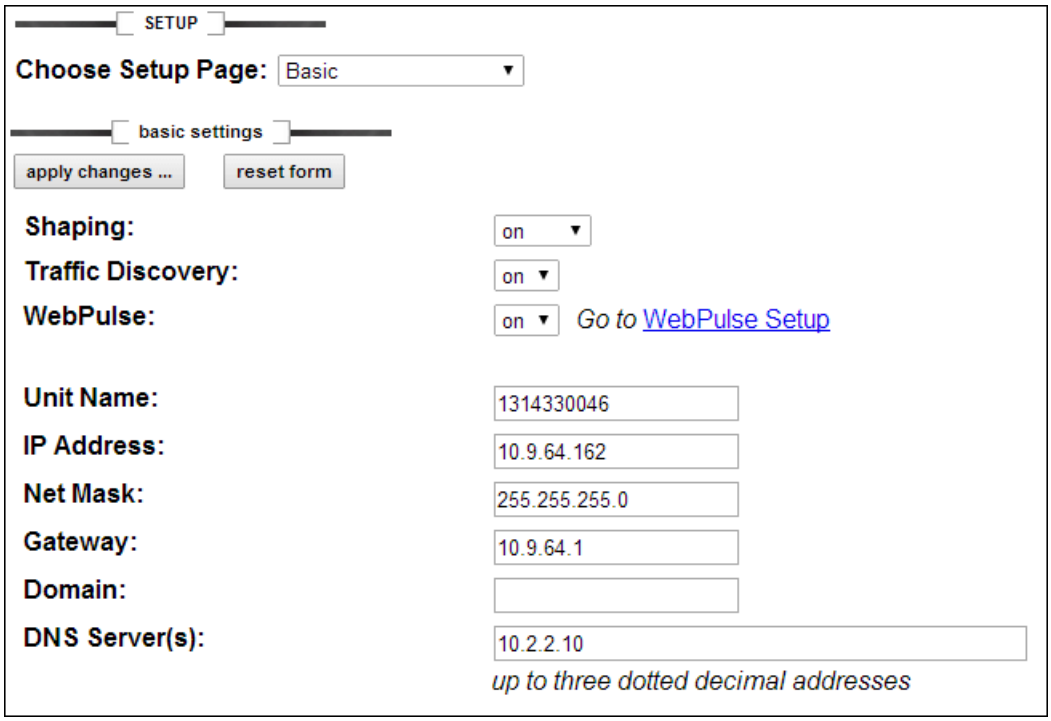

2. Verify or modify configuration details, as described in the following table.

3. Click **apply changes** to update the settings.

Field	Description
Shaping	When shaping is **on**, traffic is classified and measured, and policies and partitions actively control bandwidth allocation. When shaping is **off**, traffic is classified and measured, but it is not actively managed. When shaping is set to **watch**, PacketShaper watches the traffic on your network from a non in-line position; that is, the unit is not cabled into the main data path. Consistent with its role as a passive watcher, a PacketShaper in watch mode will not perform traffic shaping; any existing policies and partitions will be ignored. See Watch Mode Overview for more information.
Traffic Discovery	When enabled, traffic discovery lets PacketShaper automatically create classes based on the service types it detects.
URL Categories	When enabled, PacketShaper can classify by URL categories. See URL Categories Overview.
Unit Name	The identifying name or serial number of the PacketShaper. The name can be up 20 characters long and may include alphanumeric characters, dashes (-), underlines (_), and periods (.) and may not contain spaces.
IP Address	The IPv4 address assigned to the unit. The address should not be: - loopback address (127.0.0.1) - network address (all host bits 0) - broadcast address (all host bits 1) - class D or class E address - IPv6 address **Note:** You will need to reset the PacketShaper after changing its management IP address. PacketShaper will not send FDRs and link settings may not be accurate until you reset the unit.
Net Mask	The subnet mask for the network where the unit resides
Gateway	The IPv4 address the unit uses to reach other networks. The unit uses this gateway to route unit-initiated transactions to a non-local address. Specify **none** if there isn't a gateway or to clear the gateway setting.
Domain *(optional)*	The default domain name
DNS Server(s)	The IPv4 address(es) of the Domain Name Server(s). Enter a maximum of three addresses separated by spaces and/or commas. If you don't specify a DNS server, Top Talker/Listener reports will list IP addresses only, instead of more meaningful domain names.

Field	Description
Inbound Rate Outbound Rate	To manage bandwidth, PacketShaper must know the capacity of the access link it is managing. Data rates may be specified as integer bits per second, followed by a k (thousands), M (millions), G (billions), or specified symbolically (T1 , E1 , T3). ■ For a full-duplex WAN link, enter the total link speed for the inbound and outbound rates. Because full-duplex has wires that can simultaneously communicate in both inbound and outbound directions, you should enter the same rate for Inbound Rate and Outbound Rate when you expect the same amount of traffic in each direction. For example, if you have two T1 lines (3 Mbps), you should enter 3M for Inbound Rate and 3M for Outbound Rate. ■ PacketShapers can also manage a link with asymmetric Inbound and Outbound speeds, with a total rate twice the size of the link key. For example, with a 10 Mbps link key, Inbound can be configured to 18 Mbps and Outbound to 2 Mbps (for a total of 20 Mbps). This type of configuration is frequently needed on an ADSL (Asymmetric Digital Subscriber Line) link where most of the bandwidth is devoted in the downstream direction. ■ In rare situations in which the PacketShaper is managing half-duplex links, split the rate between the inbound and outbound links. For example, if you are managing a 10 Mbps half-duplex link, you could configure 5 Mbps for the inbound rate and 5 Mbps for the outbound rate since data can be transmitted in only one direction at a time. **Note:** If you have any NICs in the PacketShaper set to half-duplex, you will not be able to transmit and receive at the same time, limiting the bandwidth to half. For example, if you have set 10 Mbps half-duplex on the NIC, the PacketShaper will be able to pass only 5 Mbps IN and 5 Mbps OUT at one time. Hence, you will not be able to manage a WAN link size greater than 5 Mbps. ■ If your appliance is using NIC option cards to manage different WAN links and you don t want to control each NIC separately, the rate should be the size of the smallest NIC. For example, if you have two 1Gpbs NICs managing two links, you should specify 1Gbps for the rate. On the other hand, if you want to control each link separately, the rate should be the sum of the link speeds on all devices. For example, if Slot 1 is controlling a T1 line (1.5 Mbps) and a NIC is managing two T1 lines (3.0 Mbps), you should specify 4.5M for the rate. To control traffic across each link separately, you can create a class for each device (for example, *Slot1* and *Slot2*) and assign partitions that match the link size (1.5M for the Slot1 class and 3.0M for the Slot2 class). If your appliance is using two NICs to manage a single WAN link, specify the WAN link speed for the rate. Although the Info page will give you an error message (such as Link speed of 155 Mbps exceeds outside NIC speed of 100 Mbps) in the latter situation, it is still appropriate to specify the actual size of the link for the rate. ■ Software configuration determines maximum shaping capacity. See PacketShaper Configuration Limits.

Field	Description
	Note: 10BaseT links rarely reach the 10 Mbps limit. Keep Ethernet's practical limits in mind when configuring rates.
NIC Modes	View and modify the link speed and duplex setting for each interface. See " PacketShaper S-Series Network Interface Capabilities" on page 949 for the supported link rates and duplex modes for each available network interface. **Additional Information** ■ Some devices (such as 10 Gigabit Ethernet) support only the **auto-negotiate** setting, so the NIC mode cannot be modified. ■ If **auto-negotiate** is specified and auto-negotiation signals are not received from the other side, the negotiation will time out in one second and the interface will be set at 1000 fixed. ■ Although you can specify different fixed speeds on the Inside and Outside interfaces, such a configuration will result in a network interruption if the PacketShaper is turned off because the end devices will not be able to negotiate the correct speed for the link. ■ Whenever you wish to change Network Interface Card (NIC) settings, always select **auto-negotiate** first, then select a different value if desired. Do not change from one non-auto setting to another non-auto setting directly; re-negotiation may fail. **Caution:** When you add a new NIC, remove an existing NIC, or move a NIC to a different slot, the NIC speed and duplex settings will be reset to their defaults. After making a NIC change, make sure to verify the NIC settings and modify to desired values if necessary.

Enable/Disable Traffic Shaping

When shaping is on, traffic is classified and measured, and policies and partitions actively control bandwidth allocation. When shaping is off, traffic is classified and measured, but it is not actively managed.

> **Note:** Partitions and policies (with the exception of ignore policies) do not take effect until traffic shaping is enabled.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To turn traffic shaping on or off:

1. Click the **Setup** tab. The *Basic Settings* screen is automatically displayed in the *Setup* window. **show screen...**

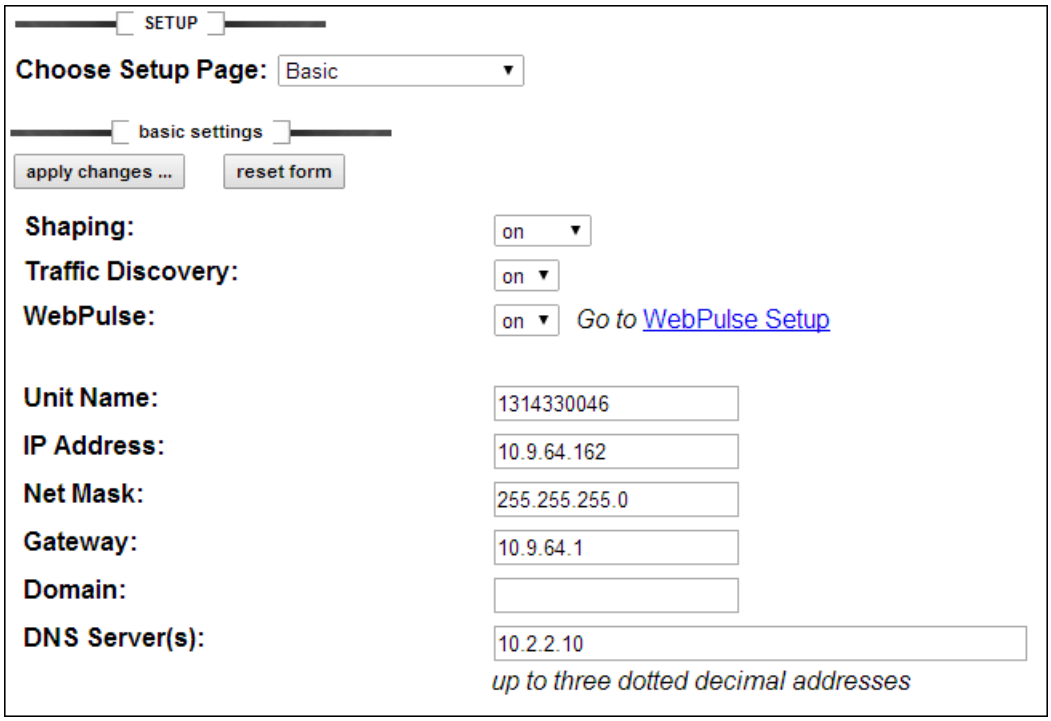

2. In the **Shaping** field, select **on** to enable shaping or **off** to disable.

3. Click **apply changes** to update the settings.

Subscribe PacketShaper Units to PolicyCenter

You can add both new and currently operating PacketShaper appliances to PolicyCenter. Units can be configured in *shared mode* or *local mode*.

- A unit running in *local mode* functions independently from other units, and has all its configuration attributes stored directly on its system disk. Once PolicyCenter is installed on a network, units in local mode can be set to shared mode and added to PolicyCenter by identifying the PolicyCenter S-Series IP address and entering its password. Once done, this unit is *subscribed* to PolicyCenter.

- A unit configured in *shared mode* can be assigned to a PolicyCenter sharable configuration, which contains settings that automatically get applied to the unit. When a unit is in shared mode, PolicyCenter continually synchronizes the unit's configuration on PolicyCenter with the configuration files on that unit's system disk; therefore, if you switch from shared mode back to local, (or the network connection to the PolicyCenter appliance is lost) the unit's configuration in local mode will be the same as its last configuration in shared mode.

When a unit subscribes to PolicyCenter, PolicyCenter creates a unique configuration for the unit at the top of the configuration tree, then assigns the unit to that configuration. The unique configuration has the unit's serial number for its configuration name, and is exclusive to the unit for which it was created—no other unit can be assigned to it unless it is renamed or moved to another location within the configuration tree. If you already have a root-level configuration named after the unit's serial number, PolicyCenter will add the unit directly to that existing configuration.

For complete documentation on PolicyCenter, refer to the PacketGuide for PolicyCenter:

https://bto.bluecoat.com/packetguide/pc1.1/index.htm

Add a New PacketShaper to PolicyCenter

To add a new, unconfigured PacketShaper to PolicyCenter, select the **shared mode** option when you run Guided Setup. When prompted, provide the IP address or host name of PolicyCenter S-Series appliance and the PolicyCenter administrator password. The PolicyCenter software will create a new PolicyCenter configuration for the unit, then automatically assign the unit to the new configuration. This unit is now subscribed to PolicyCenter.

Add a Currently Operating PacketShaper to PolicyCenter

To add a PacketShaper that is already configured on your network:

1. Access the PacketShaper's Advanced UI.
2. Click the **Setup** tab, and from the **Choose Setup Page** list, select ***PolicyCenter Access***.

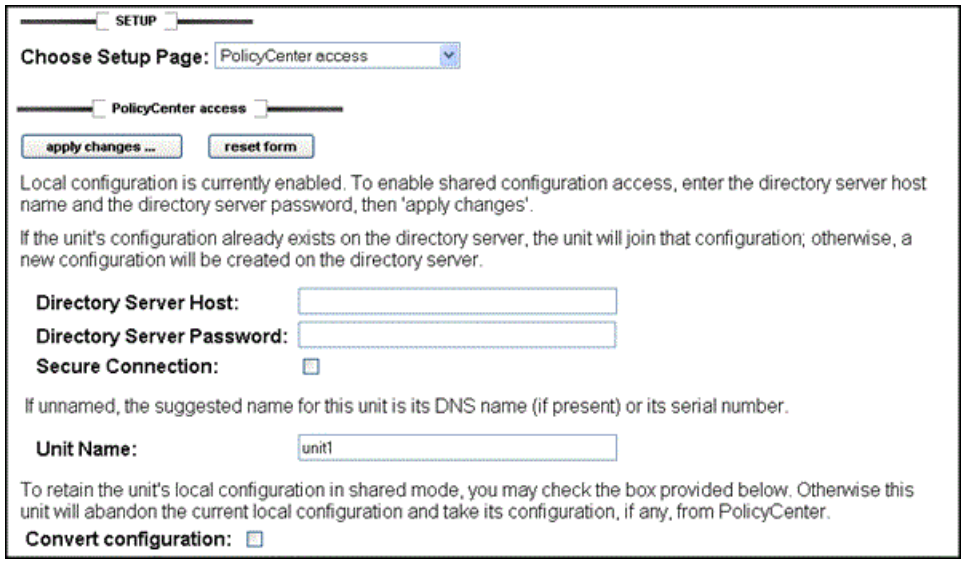

3. In the **Directory Server Host** field, provide the IP address of the PolicyCenter S-Series appliance.
4. In the **Directory Server Password** field, enter the password of your directory server. Note that this may be different from the PolicyCenter login password.
5. Check the **Secure Connection** check box to establish a secure connection between the PacketShaper and the PolicyCenter S-Series appliance.
6. In the **Unit Name** field, enter a unique name for the PacketShaper that will help you to identify the unit and its configuration within the PolicyCenter*Units* list. The suggested names are the DNS name of the unit (if present) or the unit serial number.
7. If you want a PacketShaper to retain its existing working configuration even after it has been added to PolicyCenter, select the **Convert configuration** check box. When you select this option, the unit's existing sharable attributes will be converted into a new PolicyCenter configuration with the same attributes and values. Because the unit's new PolicyCenter configuration will be based upon its previous configuration, the unit will continue to operate the same in PolicyCenter as it did in local mode. If you do not select the **convert** option, the unit's new PolicyCenter configuration is cleared, and will have default settings only.
8. Click **apply changes**. A pop-up window appears and asks you to confirm your selection.
9. Click **OK**. The unit will be prescribed to PolicyCenter, and you will be logged out of the PacketShaper browser interface.

Confirm that your PacketShaper is Installed in PolicyCenter

After you add the unit to PolicyCenter, you should confirm that PolicyCenter recognizes the new configuration.

1. Access PolicyCenter by opening another browser window, and entering the IP address of PolicyCenter.

2. Click the **Units** tab. (You may have to wait a few seconds for a response.)

If PolicyCenter displays information about your unit and its configuration, you have successfully configured your PacketShaper to work with PolicyCenter. If your unit isn't listed on the **Units** tab, refresh the PolicyCenter browser window or reopen the unit's browser interface and make sure the unit is in shared mode.

Specify Security Settings

Access to the unit can be limited in a number of ways, for example by setting passwords and selecting the protocols allowed for accessing the PacketShaper.

To view or update security settings:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **security**. The security settings appear on the *Setup* screen.

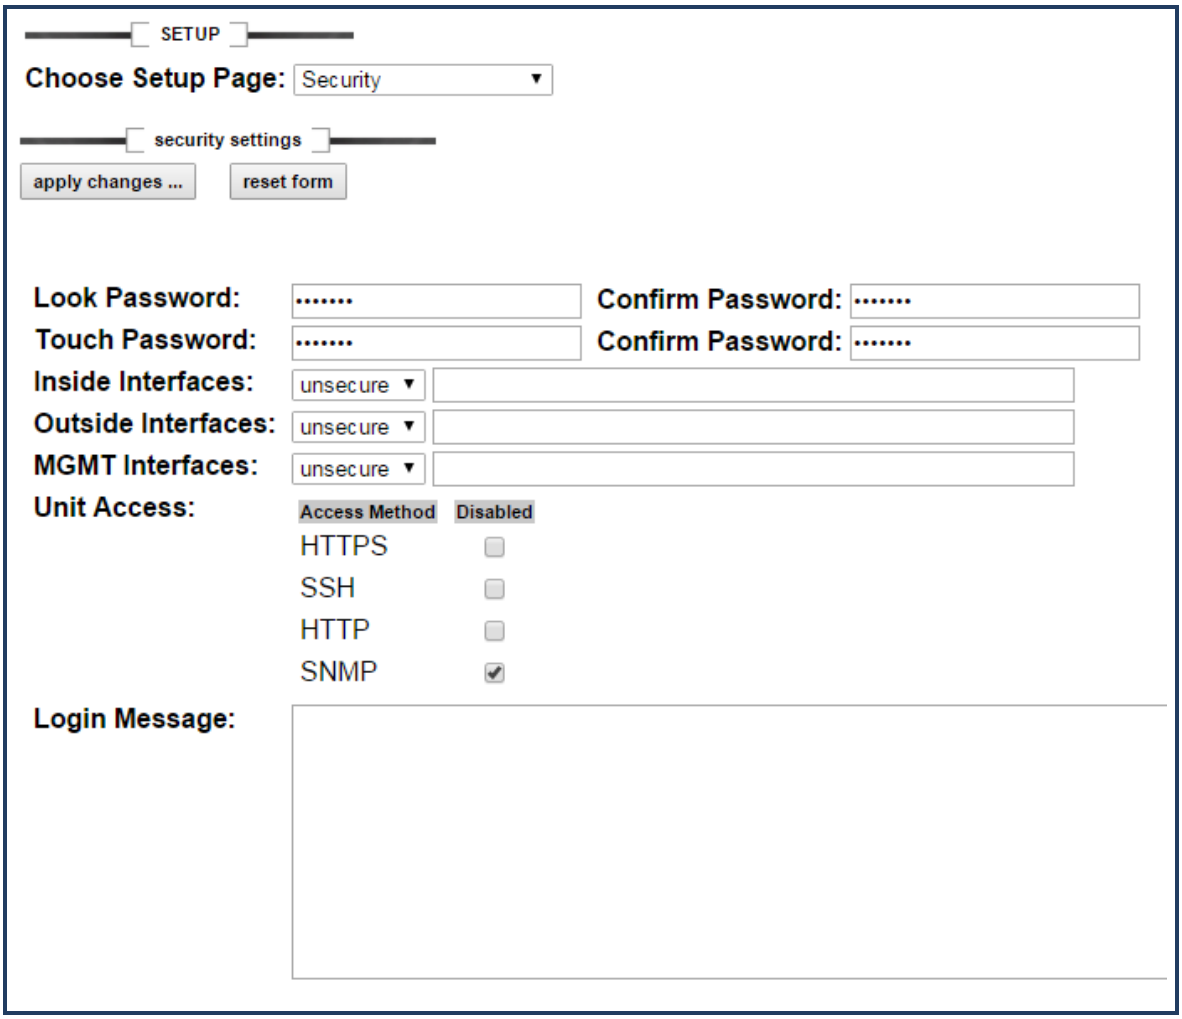

3. Verify or modify configuration details, as described in the following table.

4. Click **apply changes** to update the settings.

Field	Description
Look Password	The password for the look user that allows read-only access. To change from look to touch access, click the **log out** button on the **Info** tab, and log in as the touch user. Passwords can be up to nineteen characters long and are case-sensitive. They can consist of a combination of letters, numbers, and all special characters on the U.S. keyboard. The default password for the look user is *look*. Passwords cannot be blank. The password is set to its default when the unit is reset to factory defaults. Commands that modify the PacketShaper's configuration are not available to the look user. Similarly, you cannot retrieve sensitive information or issue commands that would impact the performance of the unit, nor can you create, edit, or delete classes, policies, or partitions. When a look user is logged in, the browser interface only presents the options that are available with look access. For example, the **class** menu is not shown on the Manage tab because look users cannot add, delete, or rename classes. You can view all the setup pages, but the **apply changes** button is not available because look users cannot change the settings. **Note:** Each time you display the Security setup page, the password fields will be populated with eight asterisks, regardless of how long the password is. These asterisks provide extra security in that they prevent anyone from discerning the length of the password string and whether a look or touch password has been set. You do not need to erase the asterisks before applying changes, unless you want to remove the password.
Touch Password	The password for the touch user that allows read/write access. After you change the touch password, you must log in again as the touch user. The touch user has access to all commands and options. Passwords can be up to nineteen characters long and are case-sensitive. They can consist of a combination of letters, numbers, and all special characters on the U.S. keyboard. The default password for the touch user is *touch*. Passwords cannot be blank. The password is set to its default when the unit is reset to factory defaults.
Confirm Password	Use this entry as verification when changing a password.

Field	Description
Inside Interfaces Outside Interfaces MGMT Interfaces	Enable/disable access to the unit over the inside, outside, or MGMT interfaces (for example, web access). When the inside, outside and MGMT interfaces are all set to **secure**, access to the unit is available only via a console connection. The browser interface is disabled. **unsecure** enables unlimited access over the specified interface. **secure** blocks all access from the specified interface. **list** enables access to up to 16 IP addresses, separated by spaces and/or commas. To specify a subnet, use the format: ipaddress:subnet_mask or ipaddress/CIDR. The list option accepts IP addresses only, not host names. To find the IP address associated with a host name, use the **dns lookup** CLI command. **Notes:** ■ Keep in mind that securing an interface means that queries such as DNS and SNTP cannot be made via the secured interface. Consider using the **list** option and including these servers and your gateway in the list. ■ The WebPulse classification features require access to a number of outside web servers. Therefore, do not completely secure the outside interface. Instead, use the **list** option and add the IP addresses of the WebPulse service points (use the **setup webpulse show service** command to find the IP addresses of the one or two fastest servers). In addition, you will need to specify the URLs of other servers PacketShaper needs to access; see "PacketShaper Port and URL Usage " on page 924for details. If you are using a web proxy, you will also need to add this server's IP address to the list if it will be accessible via the outside interface. ■ The PacketShaper will not be able to process local ARP requests via a secured interface. ■ If you secure the outside interface and your gateway is on the outside, a "gateway not found" message will be displayed in the login banner and on the info page. In this state, tasks such as upgrading the software image from a non-local address will be disabled.
Unit Access	Click the check box by a service protocol to disable access to the unit via that protocol. By default, only secure services (HTTPS and SSH) are enabled. If you are using the SNMP feature, you must enable SNMP access. **Note:** If you disable *all* protocols, you will only be able to access the unit via a direct console connection.
Login Message	Configure a message that will display before logging into the PacketShaper. The message displays on the browser login page, when logging in using a remote login utility (such as Telnet), and when console connecting to the unit. This feature is useful for informing users about the company's access policies and consequences for unauthorized use. The text can be up to 511 characters long. If you want to display a message that is longer than 511 characters, you can create a text file that contains your message text. Name the file login.txt and upload it to the 9.256/ folder. The first 2048 characters of the text file will display after any message that is configured as the Login Message. Thus, the text file is appended to the message text, allowing the message to have a total approximate length of 2500 characters.

Note: To determine whether you are logged in to PacketShaper as a look or a touch user:

Sky UI: The **Log out** link at the top indicates either **Log out (TOUCH)** or **Log out (LOOK)**.

Advanced UI: The **Info** tab indicates the Access Level underneath the **log out** button.

Password Recovery

If you forget the touch password, you can use the recovery feature to access the appliance and restore the default touch password. To ensure that security is not compromised, this feature works only when you are directly connected to the appliance's CONSOLE port.

> **Caution:** The PacketShaper is non-functional when using the recovery feature, and traffic will temporarily be blocked during the password recovery.

1. Connect the null-modem cable (included with your appliance) from the appliance's CONSOLE port to the serial port on your PC or workstation.
2. Open a terminal emulation program (such as HyperTerminal).
3. Verify that you have configured the program with the following values to communicate with the appliance's CONSOLE serial port:

   9600 bps, 8 data bits, 1 stop bit, no parity, hardware flow control
4. Power cycle the PacketShaper.
5. Log in as the recovery user:

   User name: **recovery**

   Password: **recovery**

   The recovery menu appears. **show screen...**

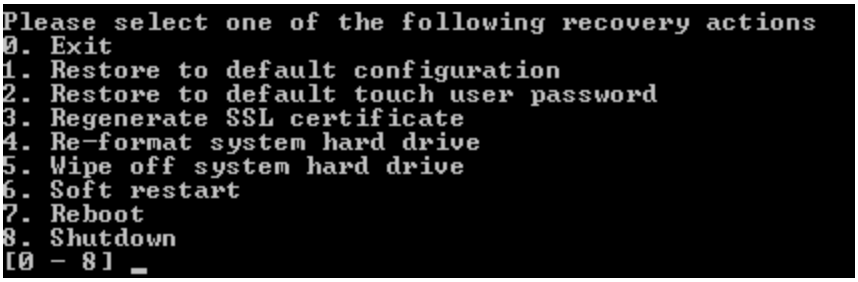

6. Type **2** to restore the default touch password.
7. Select **7** to reboot the appliance. After rebooting, the password prompt appears.
8. When prompted, enter the default touch user name (**touch**) and password (**touch**).
9. *(Recommended)* Reset the touch password. See " Specify Security Settings " on page 297.

Configure RADIUS

If you want to take advantage of your RADIUS (third-party software) user authentication, login, and accounting features, you can integrate PacketShaper and RADIUS.

Steps:

1. First, configure the RADIUS server with PacketShaper-specific attributes.
2. Configure the RADIUS authentication service in PacketShaper.
3. Configure the RADIUS accounting service in PacketShaper to have an audit trail of user logins.
4. Log into PacketShaper using RADIUS.

Configure RADIUS Servers

The PacketShaper RADIUS client has been tested with the freeRADIUS server. This section includes instructions on configuring freeRADIUS servers with PacketShaper-specific information. For information on the general setup and configuration of this and other servers, refer to the documentation included with the product.

1. Install freeRADIUS from your Linux distribution (for example, *apt-get install freeradius*).

2. Create a file named **dictionary.packeteer** (in Ubuntu, in the /usr/share/freeradius/ directory) with these lines:

```
#
 # dictionary.packeteer
 #
 # RADIUS configuration Vendor-Specific attributes by Packeteer, Inc.
 #

 VENDOR Packeteer 2334

 #
 # Standard attribute
 #

 ATTRIBUTE Packeteer-AVPair 1 string Packeteer
ATTRIBUTE Packeteer-PC-AVPair 2 string Packeteer
```

3. Add a line in the file named **dictionary** that specifies the path to the PacketShaper dictionary file (in Ubuntu, it's at /etc/freeradius/dictionary). For example:

```
$INCLUDE /usr/share/freeradius/dictionary.packeteer
```

4. In the **clients.conf** or **clients** file (in Ubuntu, it's at /etc/freeradius/clients.conf), add an entry that defines the address of the PacketShaper, and the shared secret. For example:

```
client 10.9.64.0/24 {
secret = bluecoat
}
```

5. Enter each user's name, password, and local access level (look or touch) into the **users** file. In Ubuntu, the users file is at /etc/freeradius/users. For example, for a user named Laurel who has password of *hardy* and touch access:

```
laurel Auth-Type := Local, Password == "hardy"
 Packeteer-AVPair = "access=touch"
```

Notes:

- The Auth-Type field has two possible values: Local or System. If you are using CHAP (Challenge Handshake Authentication Protocol) for your authentication method, you will use the local Auth-Type. If you are using PAP (Password Authentication Protocol), you can specify either *Local* or *System*, depending on whether the passwords are being stored and maintained by the RADIUS server (Local) or operating system (System).
- If you leave out the Packeteer-AVPair attribute, the default access level is look.
- The indentation of the Packeteer-AVPair line is required.

Configure RADIUS Logins to PolicyCenter

To use RADIUS with PolicyCenter, enter each user's name, password, and local access level (look or touch for PacketShaper) and the PolicyCenter role into the **users** file. The syntax for the value of Packeteer-PC-AVpair is *role=<org-name>:<role-name>*. For example:

```
laurel Auth-Type := Local, User-Password == "hardy"
     Packeteer-AVPair = "access=touch",
       Packeteer-PC-AVPair = "role=PC:look"
```

Configure RADIUS Authentication Service

RADIUS authentication is an optional method for users to log into the PacketShaper browser and command-line interfaces. Using third-party RADIUS servers enables you to have central configuration of user accounts.

In addition to configuring the server as described below, you need to do some configuration at the RADIUS server so that it will work with PacketShaper.

To configure PacketShaper to work with a RADIUS authentication server:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, choose **RADIUS Client**. The *RADIUS Client Settings* screen appears.

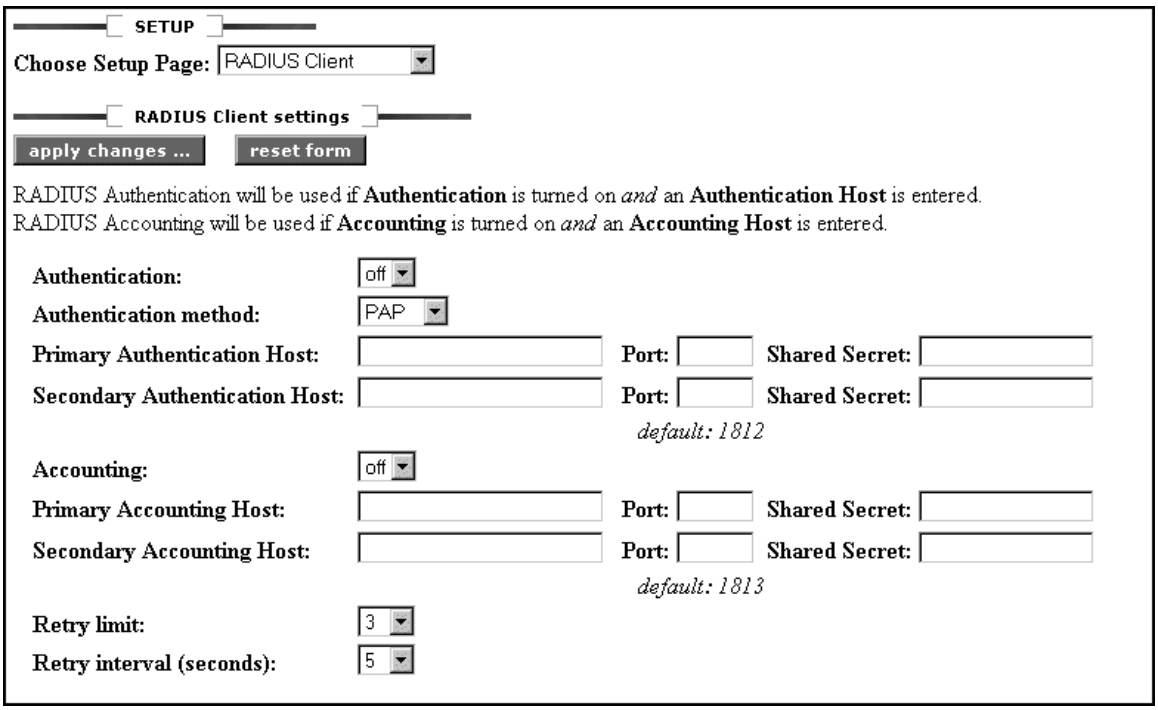

3. In the **Authentication** field, select **on**.

4. Select an **Authentication method**:

- ***PAP*** *(Password Authentication Protocol)*: With PAP, the user name and password are transmitted in clear, unencrypted text. If you select the PAP authentication method, Symantec recommends you increase security by logging into the PacketShaper browser interface via HTTPS. ASCII or PAP authentication is required for RADIUS configurations that require access to clear text passwords (for example, when passwords are stored and maintained in a database external to the RADIUS server).

- ***CHAP*** *(Challenge Handshake Authentication Protocol):* In some environments, CHAP may be preferred for greater security. The RADIUS server sends a challenge that consists of a session ID and an arbitrary challenge string, and the user name and password are encrypted before they are sent back to the server.

- ***MS-CHAP*** *(Microsoft Challenge Handshake Authentication Protocol):* This protocol is similar to CHAP, but with MS-CHAP authentication, the RADIUS server can store an encrypted version of a user password to validate the challenge response. Standard CHAP authentication requires that the server stores unencrypted passwords. If you select the MS-CHAP authentication method, Symantec recommends you increase security by logging into the PacketShaper browser interface via HTTPS.

  Note: MS-CHAP v1 and v2 are supported. PacketShaper attempts authentication with MS-CHAP v2 first. If the remote server doesn't support v2 or if authentication is denied, PacketShaper re-attempts authentication with MS-CHAP v1.

  Limitation: PacketShaper SSH does not support the MS-CHAP or CHAP authentication methods. Regardless of which method the PacketShaper is configured to use for RADIUS authentication, PAP will be used when logging in with SSH to the CLI. Note that this limitation does not apply when logging in to the web UI: the configured method will be used during authentication.

5. In the **Primary Authentication Host** field, enter the IP address or DNS name of the RADIUS server.

6. *Optional:* To access the RADIUS server with a specific port, enter a number in the **Port** field.

If the field is left blank, the default port will be used.

7. In the **Shared Secret** field, enter the designated secret.

8. *Optional:* Specify a **Secondary Authentication Host** to use in case the primary RADIUS server is not accessible. Be sure to specify its **Shared Secret** as well.

9. If necessary, adjust the **Retry limit**.

By default, if the RADIUS server fails to respond, the RADIUS client will try to log onto the server three times before reporting a server failure. You can select a value between 1 and 10. If you have specified a secondary authentication host, the RADIUS client will alternate attempts to log onto each server.

10. If necessary, adjust the **Retry interval**.

By default, the RADIUS client waits 5 seconds before retrying a login when the RADIUS server fails to respond. You can select a value between 1 and 30 seconds.

11. Click **apply changes**.

After you have configured a RADIUS authentication server, users will be prompted for a user name and password when logging into PacketShaper. For more information, see Log In and Out with RADIUS.

Configure RADIUS Accounting Service

If you want to have an audit trail for user logins, you can configure PacketShaper as a RADIUS client and specify the accounting server settings. Once this is configured, PacketShaper will send a PW_STATUS_START accounting message to the accounting server when a user logs in and a PW_STATUS_STOP message when a user logs off or is disconnected.

To configure PacketShaper to work with a RADIUS accounting server:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, choose **RADIUS Client**. The *RADIUS Client Settings* screen displays.

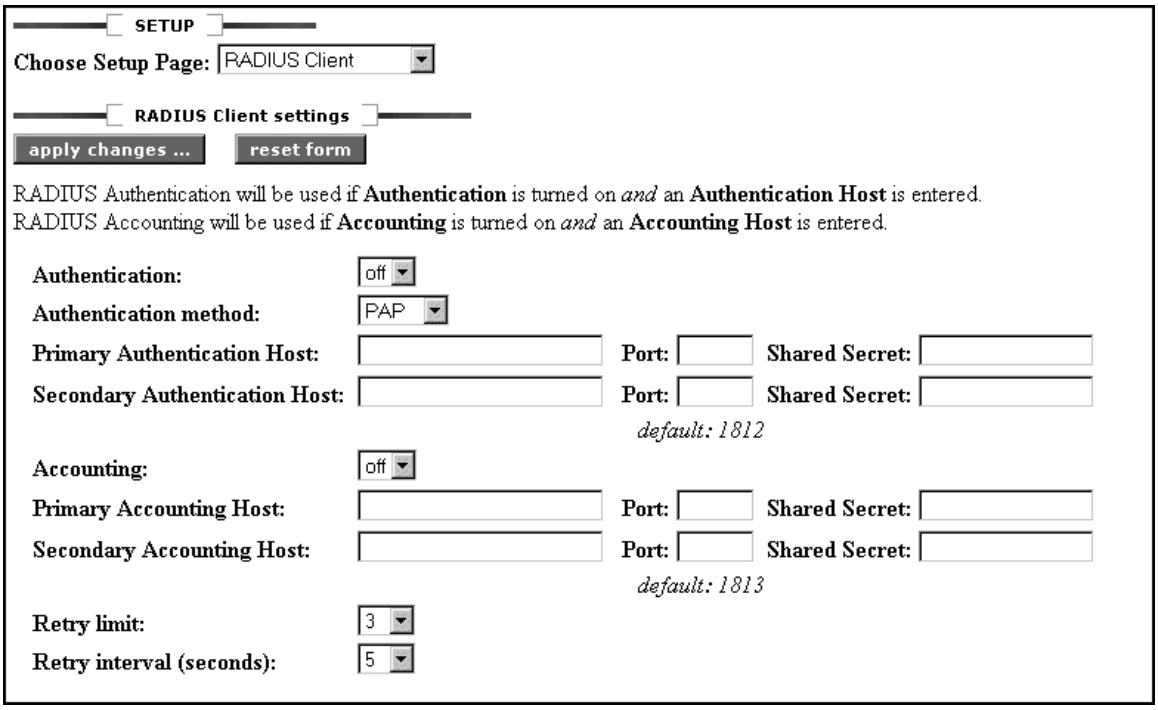

3. In the **Accounting** field, select **on**.

4. In the **Primary Accounting Host** field, enter the IP address or DNS name of the RADIUS accounting server.

5. *Optional:* To access the RADIUS server with a specific port, enter a number in the **Port** field.

If the field is left blank, the default port will be used.

6. In the **Shared Secret** field, enter the designated secret.

7. *Optional:* Specify a **Secondary Accounting Host** to use in case the primary RADIUS server is not accessible. Be sure to specify its **Shared Secret** as well.

8. If necessary, adjust the **Retry limit**.

By default, if the RADIUS server fails to respond, the RADIUS client will try to log onto the server three times before reporting a server failure. You can select a value between 1 and 10. If you have specified a secondary accounting host, the RADIUS client will alternate attempts to log onto each server.

9. If necessary, adjust the **Retry interval**.

By default, the RADIUS client waits 5 seconds before retrying a login when the RADIUS server fails to respond. You can select a value between 1 and 30 seconds.

10. Click **apply changes**.

After you have configured a RADIUS accounting server, users will be prompted for a user name and password when logging into PacketShaper. For more information, see Log In and Out.

Log In and Out with RADIUS

After the PacketShaper and the RADIUS server are configured to work together, users can log in to PacketShaper using their RADIUS credentials.

> **Caution:** CLI login using RADIUS credentials is not active until the user first logs in through the browser interface. In other words, RADIUS users must login once through the browser before they can have CLI access.

Even when RADIUS is enabled, users can still log in with their local credentials (user name of *look* or *touch*). This allows the user to log in without authenticating through the RADIUS server. This is especially useful when the RADIUS server is down or if PacketShaper is unable to connect to the RADIUS server. However, the local login technique does not record user names for auditing purposes.

Logging In with RADIUS

After RADIUS authentication and/or accounting is enabled, the user will be prompted for a user name and password when logging into the PacketShaper browser or command-line interface. The user name can be up to 63 ASCII characters and may include a realm. PacketShaper consults the configured RADIUS server to determine whether the user has access to the unit and verifies that the password is correct. The RADIUS client first tries the primary server, and if it doesn't respond within the specified retry interval, the secondary server (if configured) is tried. If there is still no response, the client repeats the primary/secondary cycle up to the specified retry limit. (Note: The retry limit and retry interval are specified on the *RADIUS Client Settings* page.)

Any failed login attempts will be sent to a Syslog server, if one has been defined.

Logging Out

For audit trail and security purposes, users should explicitly log out of PacketShaper:

- To log out of the browser interface, click the **LOG OUT** link in the banner.
- To log out of the command-line interface, type **exit**.

Logging out discards session content and generates a RADIUS accounting PW_STATUS_STOP message for the user. If a user doesn't explicitly log out, PacketShaper will automatically time out after one hour of inactivity (although the time may be learned per-session from the RADIUS server). When a PacketShaper browser session times out, a "timed out" or "unknown session" message appears the next time the user attempts to use PacketShaper. When a remote login (such as Telnet) session times out, PacketShaper sends a "timed out" message and disconnects. Note that asynchronous sessions do not time out.

> **Note:** When the Sky user interface has been loaded, the real-time graphs are constantly polling the PacketShaper so the session does not time out.

Configure TACACS+

Terminal Access Controller Access-Control System Plus (TACACS+) is a protocol that provides access control via one or more centralized servers. TACACS+ provides separate authentication, authorization and accounting services.

Steps:

1. First, configure the TACACS+ server with PacketShaper-specific attributes.
2. Configure the TACACS+ authentication service in PacketShaper.
3. Configure the TACACS+ accounting service in PacketShaper to have an audit trail of user logins.
4. Log into PacketShaper using TACACS+.

Configure Linux TACACS+ Servers using the Cisco TACACS+ Daemon

The PacketShaper TACACS+ client has been tested with the Cisco TACACS+ Daemon. This section includes instructions on configuring a Linux TACACS+ server with PacketShaper-specific information. These steps should be performed *before* you configure the TACACS+ authentication and TACACS+ accounting services via the PacketShaper browser or command-line interfaces. For more information on the general setup and configuration of your TACACS+ server, refer to the documentation included with the product.

> **Note:** This procedure only is recommended for users with previous Linux experience.

Configure the TACACS+ Server

1. Install TACACS+ server from your Linux distribution. In Ubuntu: **`apt-get install tacacs+`**
2. To configure the TACACS+ server, update the TACACS+ user configuration file *tac_plus.conf* with information for each TACACS+ user. The example configuration text below shows how to define TACACS+ users with look or touch access and a clear text password, and a touch access user with an encrypted password.

> **Note:** For additional details, refer to the *users_guide* file included with the Cisco TACACS+ files.

```
# set the secret key
 key = "<key>"

# where the accounting records should go
 accounting file = /var/log/tac_plus.acct

# users accounts
 user = <username> {
 login = cleartext "<password>"
 before authorization "echo \"access=touch\"; exit 2"
 name = "<username> touch login"
 }

user = <username> {
 login = cleartext "<password>"
 before authorization "echo \"access=look\"; exit 2"
 name = "<username> look login"
 }

user = <username> {
 login = des "<encrypt_pwd>"
 before authorization "echo \"access=touch\"; exit 2"
```

```
name = "<username> touch login"
}
```

Variable	Description
<key>	the TACACS+ secret key
<username>	User name of the TACACS+ user
<password>	Clear text password for a TACACS+ user
<encrypt_pwd>	Encrypted password for a TACACS+ user

Start the TACACS+ Server

Once your user and password settings have been configured, issue the command to start the TACACS+ server. For example:

```
/etc/init.d/tacacs_plus start
```

Configure TACACS+ Authentication Service

TACACS+ authentication is an optional method for users to log into the PacketShaper browser interface and command-line interface. Using third-party TACACS+ servers enables you to have central configuration of user accounts.

In addition to configuring the server as described below, you need to do some configuration at the TACACS+ server so that it will work with PacketShaper. (See " Configure Linux TACACS+ Servers using the Cisco TACACS+ Daemon " on page 312for an example.)

Note that the server you configure for TACACS+ authentication will also be used for authorization.

To configure PacketShaper to work with a TACACS+ authentication server:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, choose **TACACS+ Client**. The *TACACS+ Client Settings* screen appears.

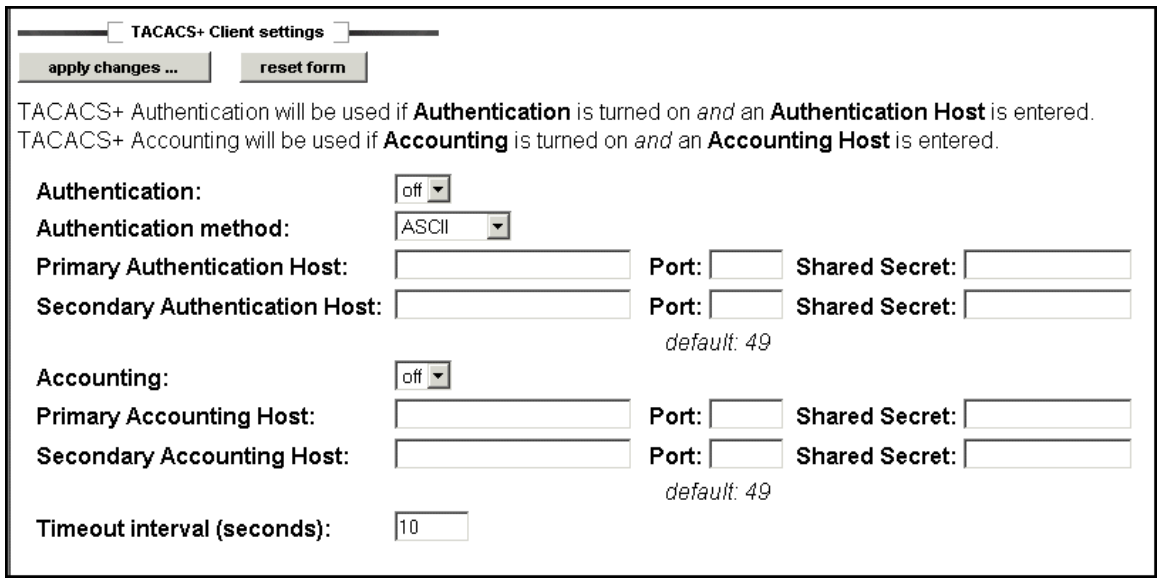

3. In the **Authentication** field, select **on**.

4. Select an **Authentication method**:

- ***ASCII*** *(American Standard Code for Information Interchange)*: With ASCII, the user name and password are transmitted in clear, unencrypted text. This is the default authentication method.

- ***PAP*** *(Password Authentication Protocol)*: With PAP, the user name and password are transmitted in clear, unencrypted text. If you select the PAP authentication method, Symantec recommends you increase security by logging into the PacketShaper browser interface via HTTPS. ASCII or PAP authentication is required for TACACS+ configurations that require access to clear text passwords (for example, when passwords are stored and maintained in a database external to the TACACS+ server).

- ***CHAP*** *(Challenge Handshake Authentication Protocol):* In other environments, CHAP may be preferred for greater security. The TACACS+ server sends a challenge that consists of a session ID and an arbitrary challenge string, and the user name and password are encrypted before they are sent back to the server.

- ***MS-CHAP*** *(Microsoft Challenge Handshake Authentication Protocol):* This protocol is similar to CHAP, but with MS-CHAP authentication, the TACACS+ server can store an encrypted version of a user password to validate the challenge response. Standard CHAP authentication requires that the server stores unencrypted passwords. If you select the MS-CHAP authentication method, Symantec recommends you increase security by logging into the PacketShaper browser interface via HTTPS.

  Notes: MS-CHAP v1 and v2 are supported. PacketShaper attempts authentication with MS-CHAP v2 first. If the remote server doesn't support v2 or if authentication is denied, PacketShaper re-attempts authentication with MS-CHAP v1. When you change the TACACS+ authentication method for a PolicyCenter sharable configuration via the PolicyCenter browser interface, TACACS+ authentication returns to the default "off" setting for that configuration. If you change the authentication method via PolicyCenter browser interface, be sure to also reenable the TACACS+ authentication feature before you apply the changes.

5. In the **Primary Authentication Host** field, enter the IP address or DNS name of the TACACS+ server.

6. *Optional:* To access the TACACS+ server with a specific port, enter a number in the **Port** field.

If the field is left blank, the default port (49) will be used.

7. In the **Shared Secret** field, enter the designated secret.

8. *Optional:* Specify a **Secondary Authentication Host** to use in case the primary TACACS+ server is not accessible. Be sure to specify its **Shared Secret** as well.

9. If necessary, adjust the **Timeout interval**.

By default, PacketShaper waits 10 seconds for a response from the TACACS+ server before the login fails. You can select a value between 1 and 60 seconds.

10. Click **apply changes**.

After you have configured a TACACS+ authentication server, users will be prompted for a user name and password when logging into PacketShaper. For more information, see Log In and Out with TACACS+.

Configure TACACS+ Accounting Service

If you want to have an audit trail for user logins, you can configure PacketShaper as a TACACS+ client and specify the accounting server settings. Once this is configured, PacketShaper will send a TAC_PLUS_ACCT_FLAG_START accounting message to the accounting server when a user logs in and a TAC_PLUS_ACCT_FLAG_STOP message when a user logs off or is disconnected.

To configure PacketShaper to work with a TACACS+ accounting server:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, choose **TACACS+ Client**. The *TACACS+ Client Settings* screen appears.

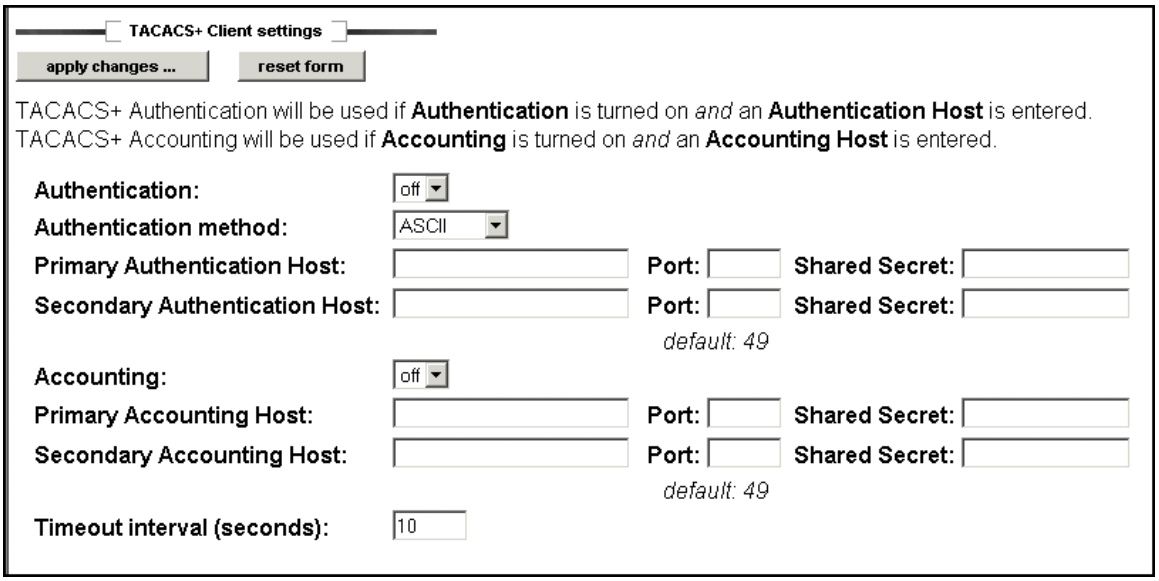

3. In the **Accounting** field, select **on**.

4. In the **Primary Accounting Host** field, enter the IP address or DNS name of the TACACS+ accounting server.

5. *Optional:* To access the TACACS+ server with a specific port, enter a number in the **Port** field.

If the field is left blank, the default port (49) will be used.

6. In the **Shared Secret** field, enter the designated secret.

7. *Optional:* Specify a **Secondary Accounting Host** to use in case the primary TACACS+ server is not accessible. Be sure to specify its **Shared Secret** as well.

8. If necessary, adjust the **Timeout interval**.

By default, PacketShaper waits 10 seconds for a response from the TACACS+ server before the login fails. You can select a value between 1 and 60 seconds.

9. Click **apply changes**.

Log In and Out with TACACS+

After the PacketShaper and the TACACS+ server are configured to work together, users can log in to PacketShaper using their TACACS+ credentials.

> **Caution:** CLI login using TACACS+ credentials is not active until the user first logs in through the browser interface. In other words, TACACS+ users must login once through the browser before they can have CLI access.

Even when TACACS+ is enabled, users can still log in with their local credentials (user name of *look* or *touch*). This allows the user to log in without authenticating through the TACACS+ server. This is especially useful when the TACACS+ server is down or if PacketShaper is unable to connect to the TACACS+ server. However, the local login technique does not record user names for auditing purposes.

Logging In with TACACS+

After TACACS+ authentication and/or accounting is enabled, the user will be prompted for a user name and password when logging into the PacketShaper browser or command-line interface. The user name can be up to 63 ASCII characters and may include a realm. PacketShaper consults the configured TACACS+ server to determine whether the user has access to the unit and verifies that the password is correct. PacketShaper first tries the primary server, and if it doesn't respond within the specified timeout interval or if the connection is refused or reset, it attempts to connect to the secondary server (if configured).

Any failed login attempts will be sent to a Syslog server, if one has been defined.

Logging Out

For audit trail and security purposes, users should explicitly log out of PacketShaper:

- To log out of the browser interface, click the **LOG OUT** link in the banner.
- To log out of the command-line interface, type **exit**.

Logging out discards session content and generates a TACACS+ accounting TAC_PLUS_ACCT_FLAG_STOP message for the user. If a user doesn't explicitly log out, PacketShaper will automatically time out after one hour of inactivity (although the time may be learned per-session from the TACACS+ server). When a PacketShaper browser session times out, a "timed out" or "unknown session" message appears the next time the user attempts to use PacketShaper. When a remote login (such as Telnet) session times out, PacketShaper sends a "timed out" message and disconnects. Note that asynchronous sessions do not time out.

> **Note:** When the Sky user interface has been loaded, the real-time graphs are constantly polling the PacketShaper so the session does not time out.

Define Flow Detail Record Collectors

PacketShaper can emit records that contain details of all flows that go through the PacketShaper to a collector on the WAN or LAN. The collector can be any product that understands Cisco NetFlow v5 flow detail records (such as Cisco Collector and Rodopi Billing Software) as well as Symantec's IntelligenceCenter. For more information, see Flow Detail Records Overview.

A collector is defined by its record type (NetFlow-5, Packeteer-1, or Packeteer-2) and its location (IP address and UDP port number). To send flow detail records to a NetFlow v5 collector, select the NetFlow-5 record type. To send flow detail records to IntelligenceCenter, select the Packeteer-2 format. The Packeteer-1 format can be emitted, but collectors for this format are not currently available.

You can define collectors with the same IP address but different record types, or with the same record type but different IP addresses. For example, you can create two collectors with the same IP address (but different ports), with one collector collecting NetFlow flow detail records and the other collecting Packeteer-2 flow detail records.

You can define up to four collectors.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To define FDR collectors:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, choose **flow detail records**. The *Flow Detail Record Settings* page appears. **show screen...**

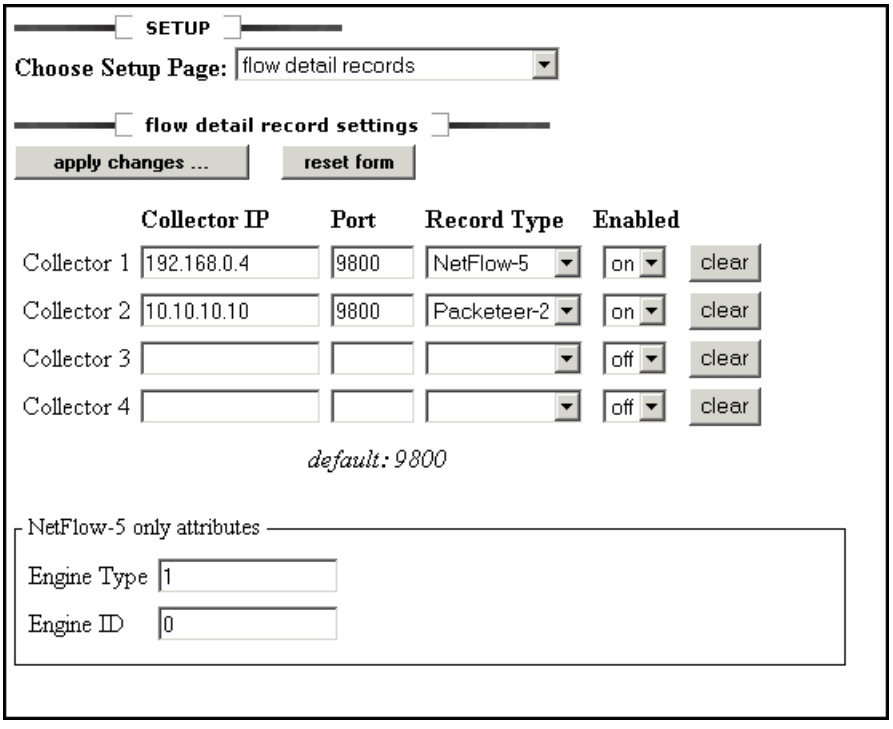

3. For each collector, specify the following:

Collector IP – IP address of the collector

Port – UDP port number of the collector (default = 9800)

Record Type – Type of record format to be emitted (**NetFlow-5**, **Packeteer-1**, or **Packeteer-2**)

Enabled – Selecting **on** enables the collector. When a collector is enabled, PacketShaper will emit flow detail records to the collector. Selecting **off** disables the collector; flow detail records will not be emitted.

> **Note:** When using the NetFlow-5 format, you can optionally assign an identifying number to a PacketShaper, in the Engine Type and Engine ID fields. You can use either, or a combination, of these fields to identify the unit that is emitting records. EngineType and EngineID are two of the fields in NetFlow-5 headers. These fields are not relevant for the Packeteer formats.

4. (*Optional for NetFlow-5 format)* Enter an integer (0-255) in the **Engine Type** and/or **Engine ID** fields.

5. Click **apply changes** to update the settings.

Enable/Disable an FDR Collector

When you want PacketShaper to begin emitting flow detail records, you enable the collector(s). You can enable a Flow Detail Records (FDR) collector as you are defining it, or you can wait until a later time.

When you want PacketShaper to stop emitting flow records, you disable the collector.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To enable or disable an FDR collector after it has been defined:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, choose **flow detail records**. The *Flow Detail Record Settings* page appears. **show screen...**

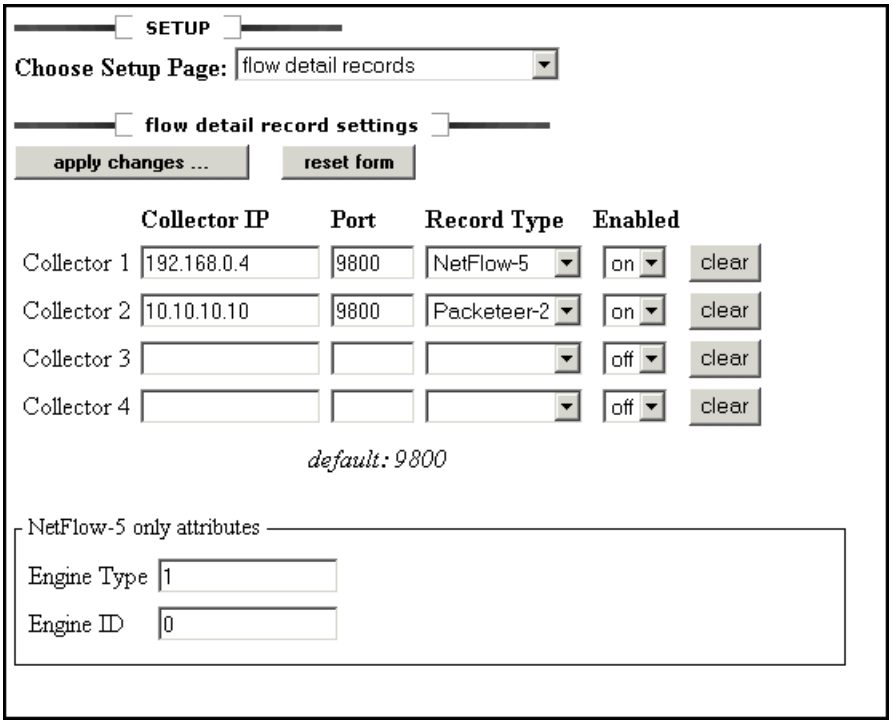

3. For each collector you want to enable, select **on** in the Enabled column. Or to disable a collector, select **off**.

When a collector is enabled, PacketShaper will emit flow detail records to the collector. When a collector is disabled, flow detail records will not be emitted, but the collector settings will be retained.

4. Click **apply changes** to update the settings.

Enable Intermediate FDRs

By default, PacketShaper emits flow detail records at the end of a flow. To change this behavior so that PacketShaper emits flow detail records at a set interval (optional), you need to enable the intermediate FDR feature.

To configure intermediate FDRs using the browser interface, do the following:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, choose **system variables**.
3. In the **Flow Detail Records (FDR)** section, change the **Intermediate FDR** setting from **off** (default) to **on**.
4. For the **Intermediate FDR Timeout** variable, specify the number of milliseconds between the emission of intermediate FDRs. The default is 15000 milliseconds.
5. To commit the changed settings, scroll to the top of the page and click **apply changes** (a unit reset is not required).

Remove an FDR Collector

If you no longer need a Flow Detail Record (FDR) collector you have previously defined, you can either disable it, define another collector in its place, or clear the settings.

PolicyCenter

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To clear an FDR collector's settings:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, choose **flow detail records**. The *Flow Detail Record Settings* page appears. **show screen...**

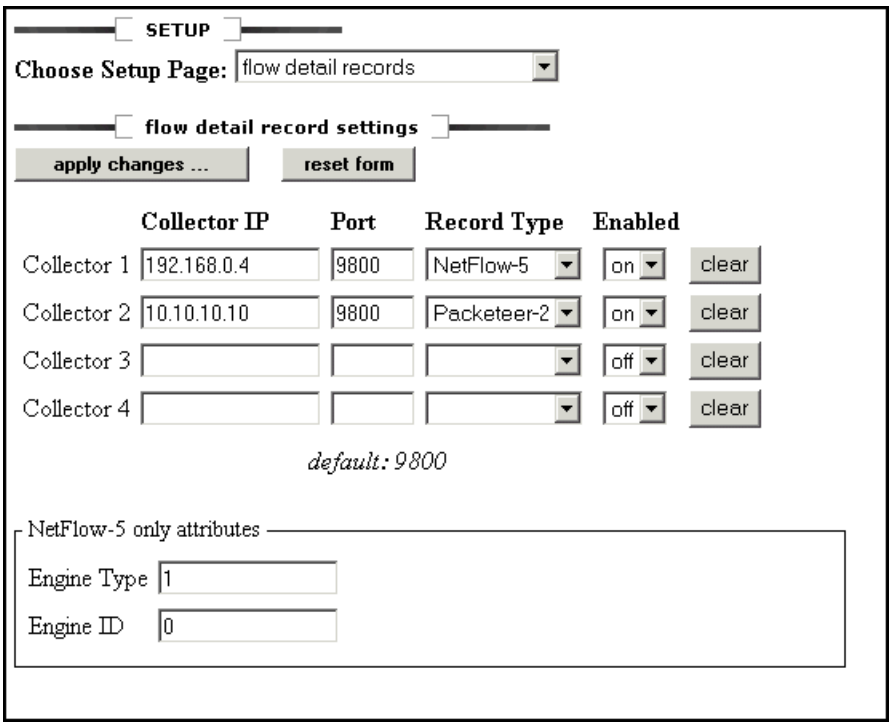

3. For each collector you want to clear, click the **clear** button. The settings for the collector are cleared.
4. Click **apply changes** to update the settings.

Modify Date and Time Settings

In addition to setting the unit's current date and time, you can select a time zone and set up time synchronization.

> **Note:** Any large-scale changing of the system time (such as changing time zones or changing the date) should always be immediately followed by a system reset so that the underlying time-sensitive scheduled operations of the PacketShaper can be correctly initialized. Changing time by a small interval should not require a reset.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To view or update date and time settings:

1. Click the **Setup** tab.
2. From the **Choose setup page** list, select **Date & Time**. The Date and Time settings appear on the Setup page. **show screen...**

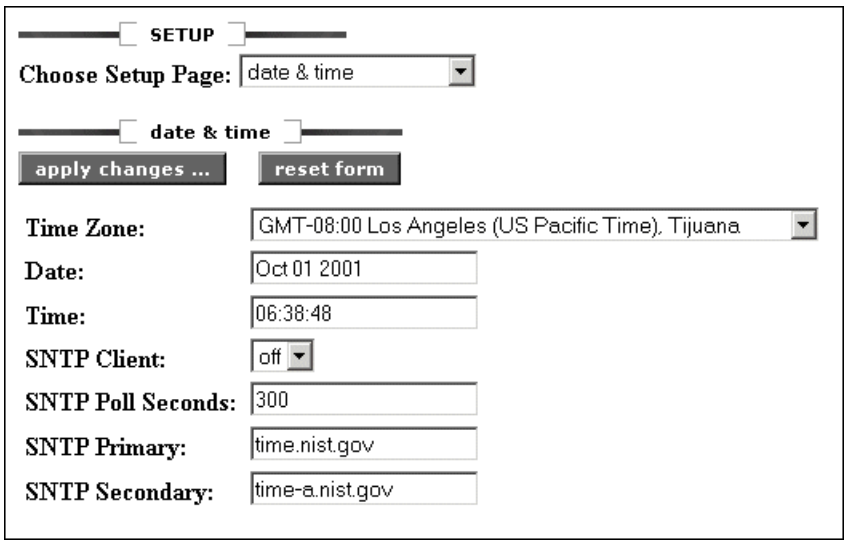

3. Verify or modify configuration details, as described in the following table.
4. Click **apply changes** to update the settings.

Field	Description
Time Zone	Select a time zone from the list. The default time zone is Pacific Standard Time, displayed as Los Angeles, CA. When you configure a time zone, PacketShaper changes its local time automatically at the start and end of daylight savings time. To configure a custom time zone, see setup timezone.

Field	Description
Date	The current date setting. Update the date by entering a new value in the format: **mmm dd yyyy** for example, *Dec 13 1999*
Time	The current time setting. Update the time by entering a new value in 24-hour format: **hh:mm:ss** for example, *20:15:00*
SNTP Client	PacketShapers are unicast Simple Network Time Protocol (SNTP) clients, enabling them to synchronize clocks with a Network Time Protocol server. Use this field to turn time synchronization on and off. Enabling SNTP is especially useful when you are managing multiple PacketShapers all clocks will be synchronized, allowing more accurate reconciliation of data when comparing reports and using a centralized reporting tool. **Note:** When time synchronization is enabled, the date and time fields become read-only. Occasionally, the default SNTP server intermittently reports that it is unavailable or that a DNS error occurred. This can happen when the default time server experiences downtime.
SNTP Poll Seconds	The number of seconds between synchronization attempts. The default is 300 seconds (5 minutes).
SNTP Primary	The DNS name or IP address of the primary time server (the default is time.nist.gov)
SNTP Secondary	The DNS name or IP address of the secondary time server (the default is time-a.nist.gov)

Upgrade the PacketShaper Software Image

Before you can add a system image to your appliance, you need to download the file from MySymantec. To download files, you must log in to your MySymantec account at the Symantec Support site.

Download the Image

To download software images, you need the following:

- The serial number of your appliance. To locate the serial number, go to the **Info** tab in the Advanced UI.
- A MySymantec account. For instructions on setting up an account, refer to the Getting Started web page.

1. Go to the Symantec Support site, MySymantec.com, and log in using your MySymantec credentials.
2. In the products list, locate the serial number of the appliance.
3. Download the .bcsi image and release notes for the desired version.
4. Place the image on a web server or workstation that the PacketShaper appliance can access.

Add a System Image

Add the image to your PacketShaper.

1. Use the File Browser to upload the image to the 9.258/image folder on your PacketShaper.

   > **Note:** Make note of the filename as you will need to specify it when adding the image in step 3 below.

2. Access the PacketShaper CLI with a remote login utility.
3. Add the image to your PacketShaper: setup images add file://*<filename>*

   For example: `setup images add file://542386.bcsi`

4. Reset the PacketShaper.
5. Clear the browser cache to see new functionality in the release.

   Firefox: **Tools > Clear Recent History > Cache**
   Internet Explorer: **Tools > Internet Options > General > Browsing History > Delete > Temporary Internet files**
   Chrome: **History > Clear browsing data > Empty the cache > Clear browsing data**

Additional Information

- To take advantage of measurement variables introduced in new versions, you must reset the measurement engine *after* installing the new version. All stored measurement data will be lost. For details, see Reset Measurement Data.
- The PacketShaper can manage up to six images. The image manager assigns each image a unique release ID (such as 532537); you must refer to this ID when activating, removing, locking, and unlocking images. If your PacketShaper already has six images installed and you add another image, the oldest unlocked image will be replaced with the new image.
- If an image fails to boot, the image manager marks this image as corrupted and then loads the last image that successfully booted.
- To see which image is the active one, use the version command. The release ID of the active image appears next to *build*.
- Image files have a .bcsi file extension. Symantec recommends that you not rename the extension portion of the filename.

Recovery Feature

The recovery feature includes the following utilities:

- **Restore to default configuration** If the PacketShaper gets into a state where it is continuously rebooting and fails to present the login prompt, you can use the recovery feature to regain access to the appliance. This utility clears all settings and restores them to the factory defaults, so it is only used in situations when you cannot otherwise log in to the PacketShaper. You will subsequently need to reconfigure the appliance.
- **Restore to default touch password** This utility is useful if you are unable to log in because you have forgotten the touch password. (See "Password Recovery" on page 301.)
- **Regenerate SSL certificate** If you believe the PacketShaper's self-signed certificate's security was compromised, you can generate a new digital certificate.
- **Re-format system hard drive** Removes the hard drive partitions on the active drive. This option is a quick way to fix a hard drive that has become corrupted. It should only be used under the direction of a Symantec Support representative.
- **Wipe off system hard drive** Erases the selected hard drive and fills the drive with zeros. Use this option, under the direction of a Symantec Support representative, if you need to swap the hard drives or before returning the drive with an RMA.

To use the recovery feature, you must directly connect to the PacketShaper CONSOLE port and log in as the recovery user. After you are logged in, the recovery menu displays.

> **Caution:** The PacketShaper is non-functional when using the recovery feature.

To ensure that security is not compromised, the recovery feature works only when you are directly connected to the appliance's CONSOLE port.

1. Connect the null-modem cable (included with your appliance) from the appliance's CONSOLE port to the serial port on your PC or workstation.
2. Open a terminal emulation program (such as HyperTerminal).
3. Verify that you have configured the program with the following values to communicate with the appliance's CONSOLE serial port:

   9600 bps, 8 data bits, 1 stop bit, no parity, hardware flow control
4. Power cycle the PacketShaper.
5. Log in as the recovery user:

   User name: **recovery**

   Password: **recovery**

The recovery menu appears. **show screen...**

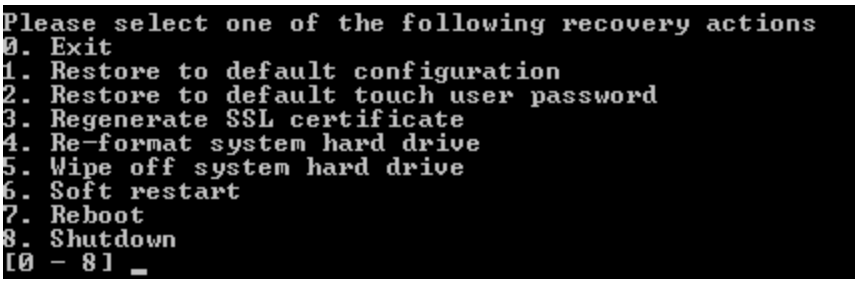

6. Enter the number that corresponds to the utility you want to run. For example, type **1** and press Enter to restore to default configuration.
7. After the utility is finished, select **7** to reboot the appliance. After rebooting, the password prompt appears.
8. When prompted, enter the default touch user name (**touch**) and password (**touch**).

   If you restored the default configuration or reformatted/wiped the hard drive, the following message appears and you will be asked if you want to run the Guided Setup wizard:

   ```
   PacketShaper v11.9 2015-08-06
   Copyright (c) 1996-2015 Blue Coat Systems, Inc. All rights reserved.
    PacketShaper not yet configured. Do you wish to be
    guided through initial setup of the PacketShaper? (yes):
   ```

9. Press Enter to run Guided Setup. You will need to reconfigure the PacketShaper IP address and other basic settings. See "Pre-Installation Check List " on page 1104.
10. If there was an issue with the image you were running before the recovery, you may want to activate an older, more stable image. Use the setup images show command to see a list of images and their release IDs, and then use setup images active to select the image to activate.
11. *(Optional)* If you have a recent backup of the configuration, you can restore it. See " Restore a Configuration" on page 710.
12. *(Recommended)* Reset the look and touch passwords. See " Specify Security Settings " on page 297.

Configure an SMTP Server

Before PacketShaper can send out email messages, you must configure an SMTP (Simple Mail Transfer Protocol) server.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To configure an email server:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **email**. The email settings appear on the *Setup* screen.

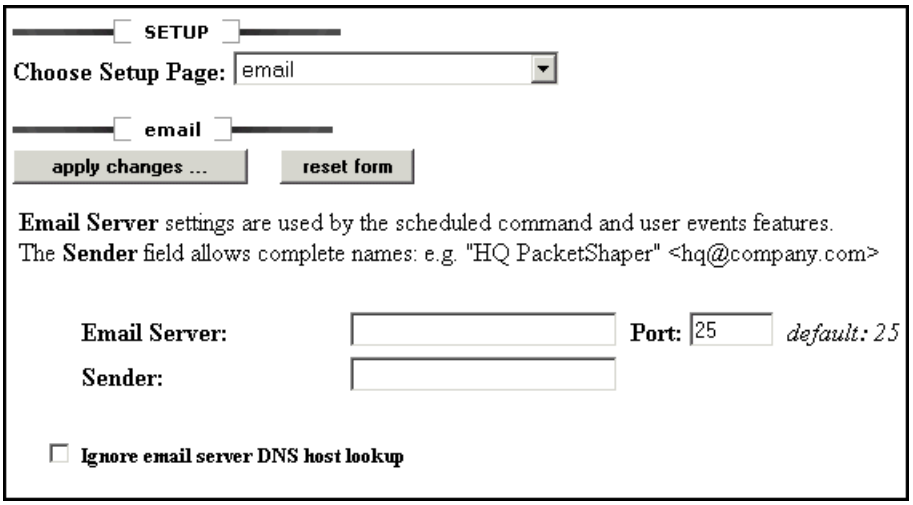

3. In the **Email Server** field, specify the SMTP server using either its DNS name or IP address.

4. To specify a non-standard port for email messages, enter the port number in the **Port** field. The default SMTP port is 25.

5. *(Optional)* In the **Sender** field, specify a complete mail address, including the domain name for example, *john_doe@example.com.* The sender's name will appear in the From line of any email message that the user event or scheduled command feature sends out.

The full form can also be used for the sender that is, a quoted name followed by the explicit email address. For example: *"Bob" <bob@examplecompany.com>*

6. If you are defining the email server before it is accessible, select the **Ignore email server DNS host lookup** check box.

> **Note:** If you don't select the check box and the email server doesn't exist yet, PacketShaper will do a DNS lookup of the email server, and when it can't find it, will display an error message. The email settings will not be accepted.

7. Click **apply changes** to update the settings.

Reset the Traffic Tree

If you want to restore the traffic tree to one of the factory-default configurations, you can reset the tree. This operation clears all classes, policies, and partitions and reverts to one of the following trees:

- **Default**–This bare-bones traffic tree includes: /Inbound and /Outbound classes with a Default class for each, and a Localhost class for the inbound and outbound directions. You can build out this tree by turning on traffic discovery or by manually creating classes for the type of traffic you want to track.

- **Model**–A pre-configured tree that can be used as is, or modified to suit your needs. This tree organizes network traffic into folders of typical genres, such as VoIP, risky websites, business-critical applications and data, and recreational websites and applications. It includes classes based on service groups and URL categories.

To reset the traffic tree:

1. Use the config save CLI command to back up your configuration before resetting the tree; this gives you the capability of restoring the traffic tree and configuration if you change your mind.

2. Click the **Setup** tab.

3. From the **Choose Setup Page** list, choose **unit resets**. The Unit Resets options appear on the *Setup* screen. **show screen...**

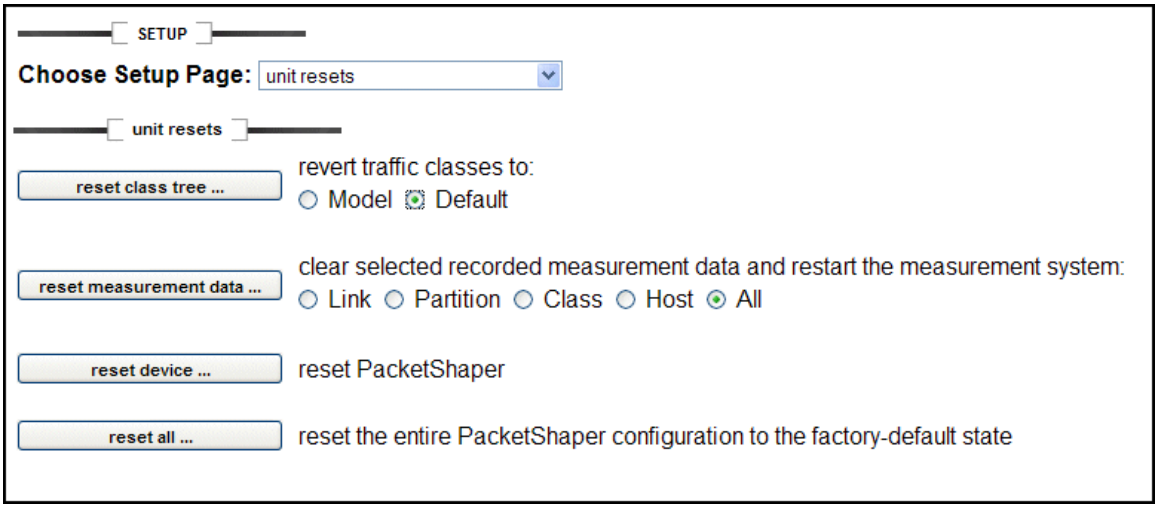

4. Click **Model** or **Default**.

5. Click **reset class tree**.

6. Click **OK** to confirm.

Additional Information

- When you reset the class tree to the Model, the following global settings will take effect: Traffic Discovery will be disabled and URL Categories will be enabled.

- If a number of classes have top talkers/listeners enabled, resetting the traffic tree could take several minutes when many hosts are being tracked. Before resetting the class tree, you may want to disable top talkers/listeners to avoid this delay.

Reboot the Unit

Resetting/rebooting the appliance performs a proper UNIX reboot, without losing power to the sytstem. In certain situations, you may need to reboot the unit. For example, you will need to reboot after after changing some system variables.

To reboot the unit:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, choose **unit resets**. The Unit Resets options appear on the *Setup* screen. **show screen...**

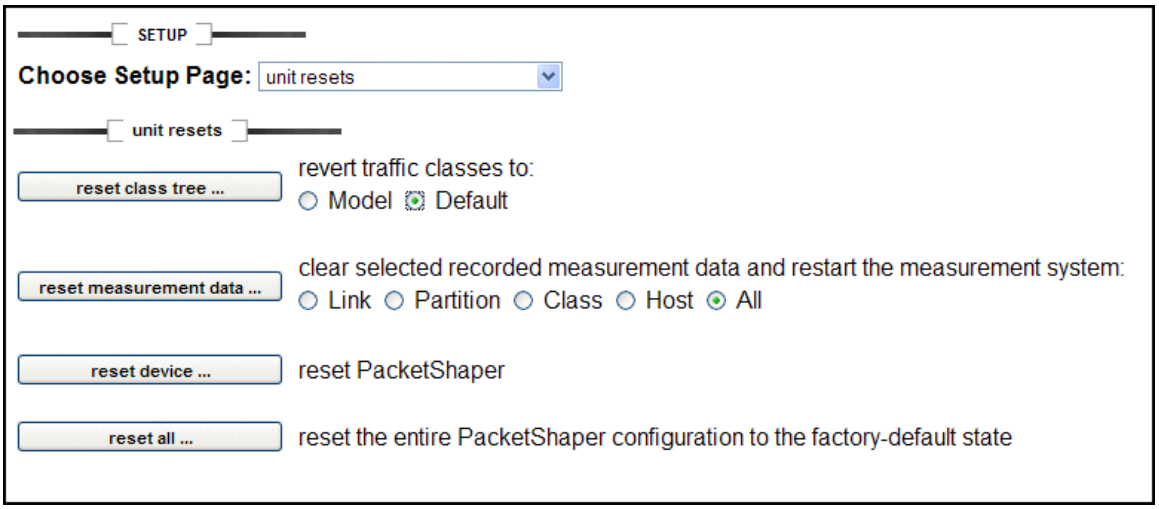

3. Click **reset device**.

The statistics that are cleared after you reset the device include: active IP hosts, current speed per host, active flows, peak flows, Top Talkers/Listeners,. All counters on the Monitor Traffic page are reset. All traffic configurations (such as classes, policies, and partitions) are preserved when you reboot.

Reset Unit to Factory Default Settings

The quickest way to clear all the PacketShaper configuration settings is to use the **Reset All** command. This command resets *all* configurations to the factory default settings:

- All classes are removed, and the traffic tree reverts to its original configuration: /Inbound and /Outbound folders with a Default class for each and a Localhost class for the inbound and outbound directions.
- All policies and partitions are removed.
- All measurement data is cleared.
- Passwords are reset to their defaults (*look* for the look user, *touch* for the touch user).

To reset the appliance:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, choose **unit resets**. The Unit Resets options appear on the *Setup* screen. **show screen...**

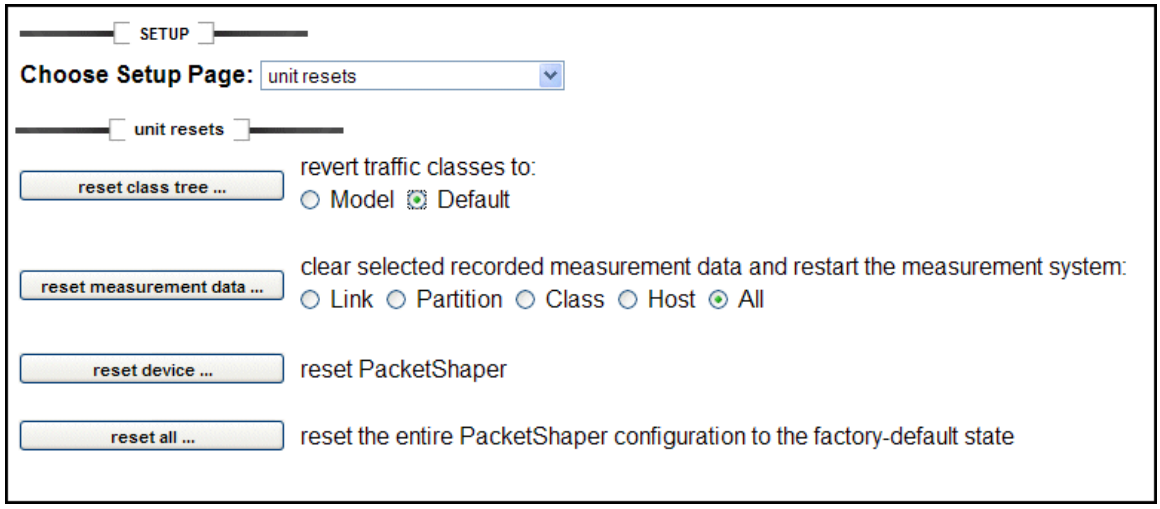

3. Click **reset all**.

This action removes the appliance's IP address, making it unreachable on your network until it is reconfigured using Guided Setup. See Run Guided Setup for details.

> **Note:** Another way to reset the appliance to factory default settings is with the recovery feature. You would use the recovery feature if you were unable to log in to the appliance for some reason (for example, if the appliance was continuously rebooting or if you forgot the touch password).

Configure SNMP on PacketShaper

The Simple Network Management Protocol (SNMP) is a widely used method of monitoring computer networks. You can configure PacketShaper to automatically send messages, such as warnings of alarm conditions (for example, Link Down), to any SNMP trap listener.

> **Note:** You must enable SNMP access to use the SNMP feature on PacketShaper. See " Specify Security Settings " on page 297.

> **Tip:** Since each individual SNMP setting change requires a full configuration reload, Symantec recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To configure SNMP on PacketShaper:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, select **SNMP**.

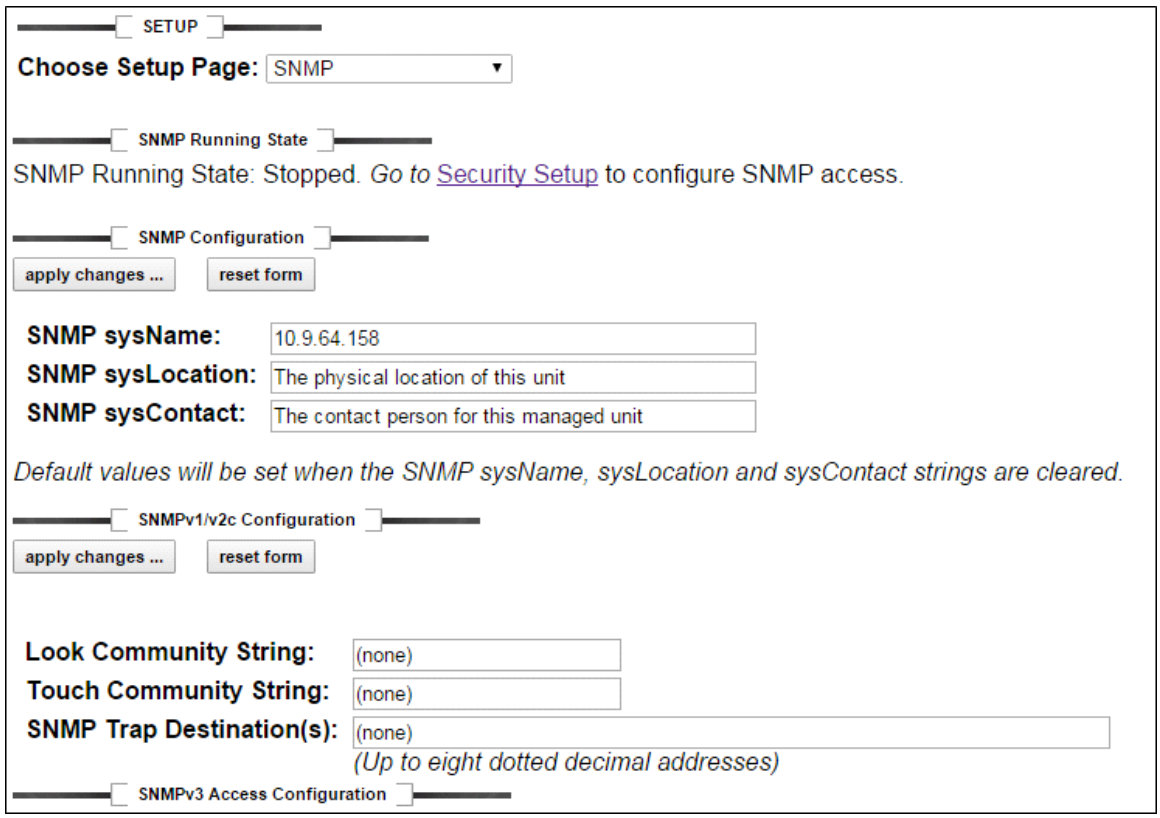

3. The **SNMP Configuration** section of this screen displays the PacketShaper's SNMP *sysName*. The default value for this setting is the unit's IP address, but you can change this setting by entering the unit's fully-qualified domain name in the **SNMP sysName** field.
4. In the **SNMP sysLocation** field, enter the physical location of the PacketShaper (room, floor, building). The location name can be up to 256 characters long.
5. In the **SNMP sysContact** field, enter the name of the person managing the unit. The contact name can be up to 256 characters long.
6. Click **apply changes**.

Configure PacketShaper for SNMPv1-v2c Support

In addition to configuring basic SNMP settings, you must also configure settings specific to SNMPv1-v2c:

- *Community strings*. Similar to a password. You must set the look community string so that PacketShaper MIB information can be accessed by network management platforms and network management stations.
- *Trap destinations*. An IP address of a system that receives SNMP traps (a *trap listener*). You can define up to eight trap destinations.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

1. On the SNMP setup page, go to the **SNMPv1/v2c configuration** section. **show screen...**

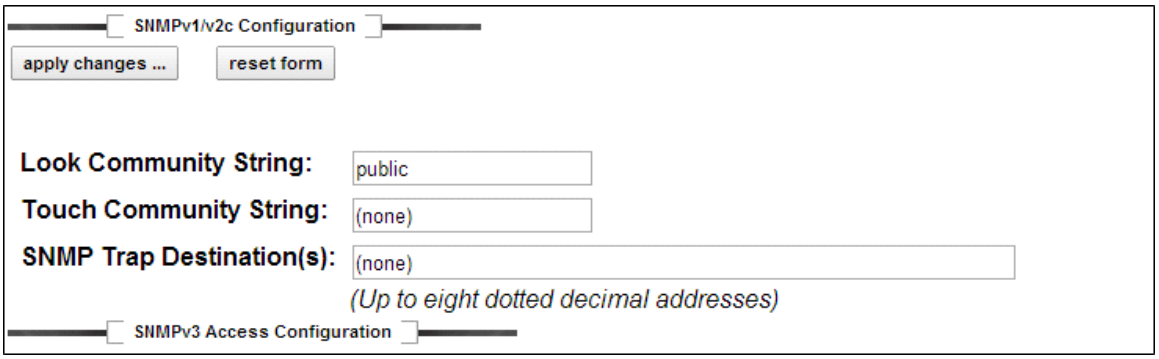

2. Set the look (read) and touch (read/write) community strings used by SNMP for authentication. By default, the SNMP Look community string is set to *public*. Allowed characters are: [A-Z] [a-z] [0-9] [_-.,@!$%^*~=+]
3. In the **SNMP Trap Destination(s)** field, enter up to eight SNMP trap destinations, specified as IP addresses (*not* DNS names), separated by spaces.

> **Note:** If you enter an IP address in the **SNMP Trap Destination** field, PacketShaper sends traps with the community string *public*. If you want PacketShaper to send the trap with a private community string, do not enter the IP address in the **SNMP Trap Destination** field; instead use the **Notify Targets** option under SNMPv3. The **Notify Targets** option can send SNMP v1/v2c/v3 traps with the custom Community Name you configure. See Configure the Notify Targets Table.

4. Click **apply changes** .
5. Make sure SNMP access is enabled. See " Specify Security Settings " on page 297.

Configure PacketShaper for SNMPv3 Support

In addition to configuring basic SNMP settings on PacketShaper, you can define SNMPv3-specific settings: views, access groups, users, and notify targets.

> **Note:** You must enable SNMP access to use the SNMP feature on PacketShaper. See " Specify Security Settings " on page 297.

> **Tip:** Since each individual SNMP setting change requires a full configuration reload, Symantec recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

Configure SNMP Access Tables

To allow or limit SNMP users' access to MIB information, you can define entries in each of the SNMP Access Configuration tables. Click on any of the links below for details on configuring the SNMPv3 tables. Note that each SNMPv3 table supports up to 32 separate entries.

- *Views* Table
- *Access Groups* Table
- *Users* Table

Configure the *Views* Table

An SNMP view filters objects from the entire MIB and defines a subset of MIB objects. Every SNMP access group has views for read and write access which either allow or limit that group's access to MIB objects. There are two predefined views; *isoAll* and *isoNone*. The *isoAll* view gives a group access to all MIB information, and the *isoNone* view blocks all access.

If you want your SNMP groups to have either complete access or no access to all MIB information, your groups only need to use the built-in *isoAll* or *isoNone* views. If, however, you want a group to access just a subset of MIB information, you will have to create a new view that describes those MIB object identifiers (OIDs) that should be included or excluded.

Create a New View

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To define a new SNMP view:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, select **SNMP.**
3. Go to the **SNMPv3 Access Configuration** section.
4. Click the **new**button beside the **Views**heading. The*Add a New View*window opens. **show screen...**

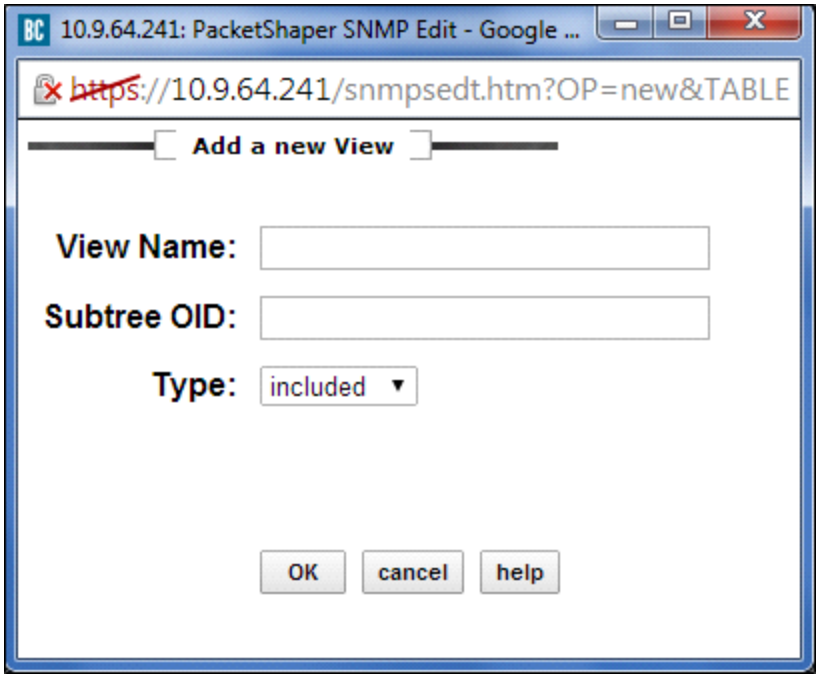

5. In the **View Name** field, enter a name for the new view. A view name can have up to 32 characters; hyphens, underscores, and periods are acceptable. Each SNMP view name must be unique.
6. In the **Subtree OID** field, enter an OID name, number, or an initial OID name and a number, such as *packeteerMibs, 1.3.6.1.4.1.2334.2,* or *packeteerMibs.1.4.* This parameter also supports the use of asterisks as wildcards for OID numbers, for example, *interfaces.*.*.1*

> **Note:** The *Add a New View* window only allows you to specify a single OID. You can, however, add additional OIDs to an existing view.

7. Click the **Type** drop-down list, and select **included** if the OID name or number should be included in this view, or **excluded** if it should be excluded from the view.
8. Click **OK** to save your settings. The new view will appear in the *Views* table on the *SNMP* setup screen.

Add Additional OIDs to a View

When you first create a view, you can only specify a single OID name or number for that view. To add additional OIDs to an existing view:

1. If you are not already on the *SNMP* setup page, click the **Setup** tab, and select **SNMP** from the **Choose Setup Page** list.
2. Click the**add oid** button beside the*Views*table entry you want to edit. The*Add an OID to View* window opens.**show screen...**

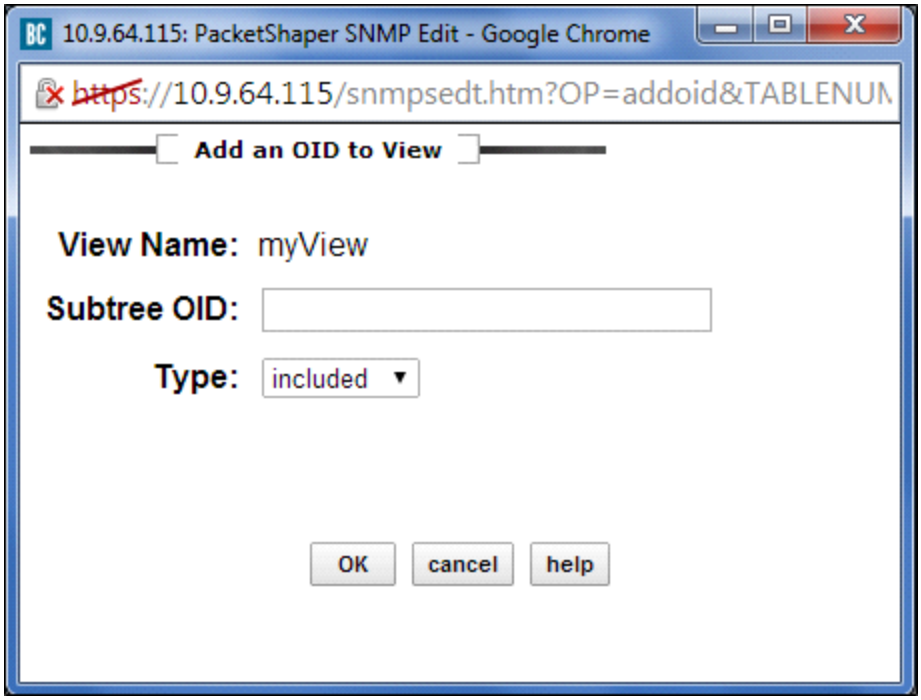

3. In the **Subtree OID** field, enter an OID name, number, or an initial OID name and a number, such as *packeteerMibs, 1.3.6.1.4.1.2334.2,* or *packeteerMibs.1.4.* This parameter also supports the use of asterisks as wildcards for OID numbers, for example, *interfaces.*.*.1.*
4. Click **OK** to save your settings. The updated view will appear in the *Views* table on the *SNMP* setup screen.

Edit an SNMP View

Change the view type from **included** to **exclude**d, or replace a view's existing OID(s) with a new OID.

1. If you are not already on the *SNMP* setup page, click the **Setup** tab, and select **SNMP** from the **Choose Setup Page** list.
2. Click the **edit** button beside the *Views* table entry you want to edit. The*Edit View*window opens.**show screen...**

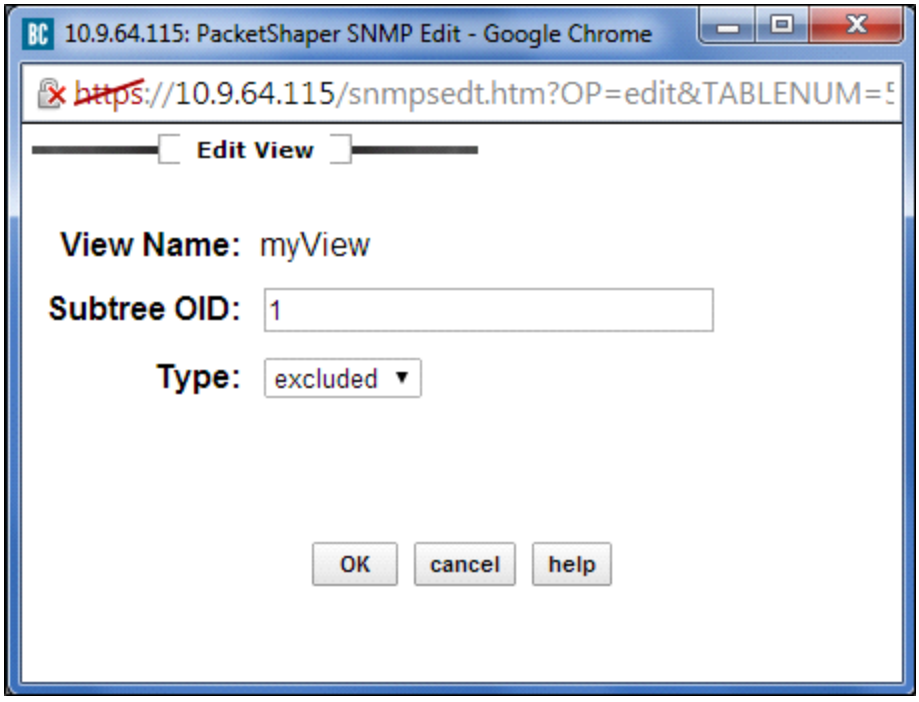

3. Update the view's). **Subtree OID**or**type**settings (described in Create a New View**Note**: When you edit an SNMP view, all its previously configured OIDs are cleared, and are replaced by the single OID in this field. Any additional OIDs for this view must be manually added again using the procedure described above.
4. Click **OK** to save your settings. The updated view will appear in the *Views* table on the **SNMP** setup screen.

Delete an SNMP View

Before you delete an unwanted SNMP view, verify that no access groups or targets are using that view. You will not be able to delete a view referenced by other SNMP tables. To delete an SNMP view:

1. If you are not already on the *SNMP* setup page, click the **Setup tab**, and select **SNMP** from the **Choose Setup Page** list.
2. Click the **delete** button beside the view you want to remove. A popup window asks you to verify that the view should be deleted.
3. Click **OK**.

Configure the *Access Groups* Table

Each SNMP access group is defined by a group name, a security model (and level), and a set of views that specifies which types of MIB data that access group can read or write.

Create a New Access Group

PolicyCenter

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a new access group:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, select **SNMP.**
3. Click the**new**button beside the**Access Groups** heading. The*Add a New Access Group* window opens. **show screen...**

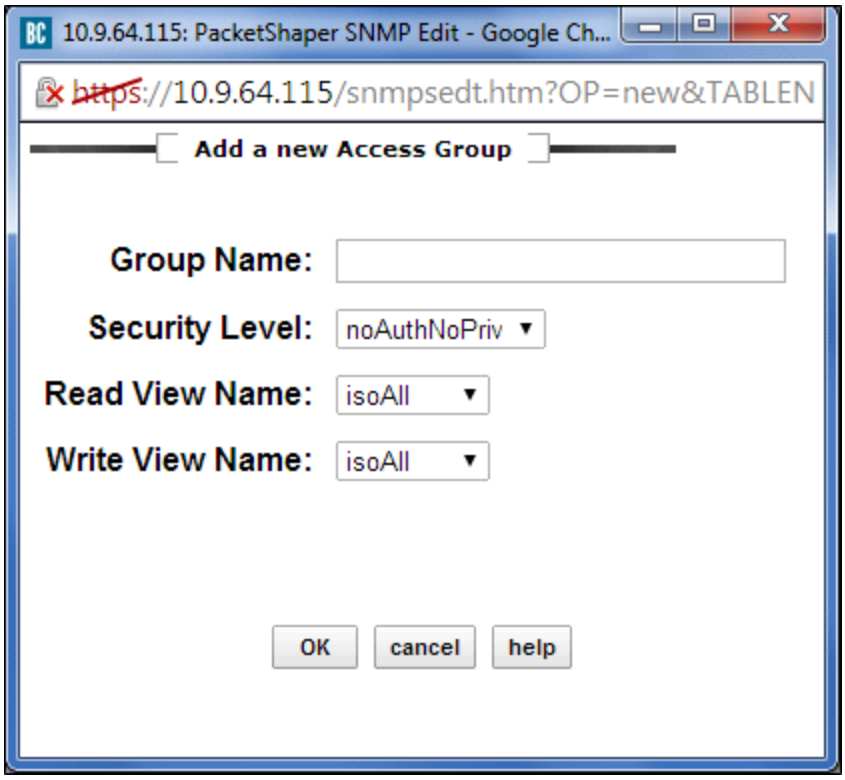

4. In the **Group Name** field, enter the name of the new access group. An access group name can be up to 32 characters; hyphens, underscores, and periods are acceptable. Each access group name must be unique.

5. Click the **Security Level** drop-down list, and select one of the following security levels for that group.
   - **noAuthNoPriv**: Identifies a user for access control, but does not provide authentication.
   - **authNoPriv**: Identifies a user for access control, and authenticates the user's password.
   - **authPriv**: Identifies a user for access control, authenticates the user's password, and provides encryption.
6. Access groups have *read* (look) access to the information specified by the read view. Click the **Read View Name** drop-down list and select the predefined view **isoAll** (to allow complete read access) or **isoNone** (to block all access), or click the name of a user-defined custom view.
7. Access groups have *write* (touch) access to the information specified by the write view. Click the **Write ViewName** drop-down list and select the predefined view **isoAll** (to allow complete write access) or **isoNone** (to block all access), or click the name of a user-defined custom view.

8. Click **OK** to save your settings. The new access group will appear in the *Access Groups* table on the *SNMP* setup screen.

Edit an Access Group

Modify security models or levels, or change a group's read and write views. To edit an access group:

1. If you are not already on the *SNMP* setup page, click the **Setup** tab, and select **SNMP** from the **Choose Setup Page** list.
2. Click the**edit** button beside the*Access Groups* table entry you want to edit. The*Edit Access Group* window opens. **show screen...**

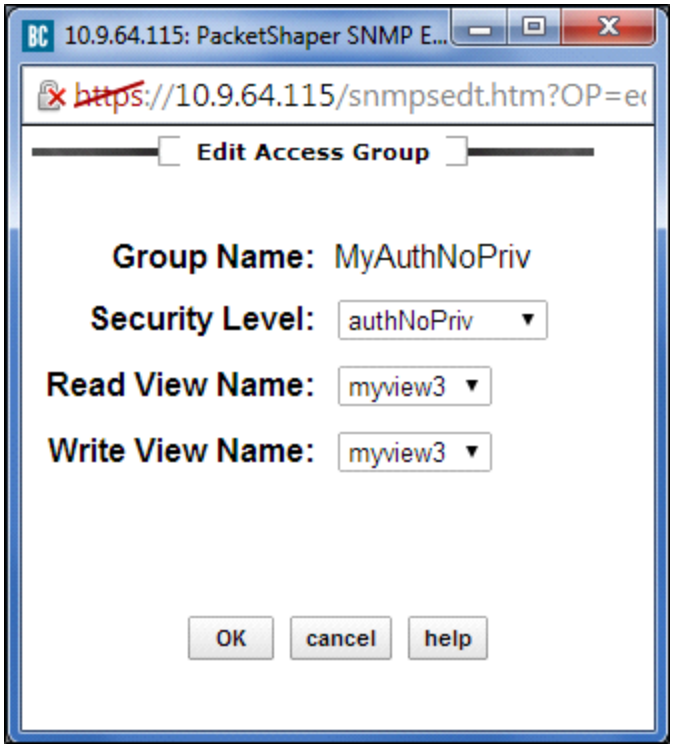

3. Update the access group's security and view settings (described in Create a new Access Group).
4. Click **OK** to save your settings. The updated access group will appear in the *Access Groups* table on the *SNMP* setup screen.

Delete an Access Group

You cannot delete an access group that is being used by any SNMP user. Before you attempt to delete an access group, check the **Group Name** column of the *Users* table, and verify that no user is referencing the access group you want to delete. Note that you can edit SNMP users and reassign them to different access groups if necessary.

To delete an access group from the *Access Groups* table:

1. If you are not already on the *SNMP* setup page, click the **Setup tab**, and select **SNMP** from the **Choose Setup Page** list.

2. Click the **delete** button beside the group you want to remove. A popup window asks you to verify that the access group should be deleted.
3. Click **OK**.

Configure the *Users* Table

Each SNMP user entry defines a user (login) name, an association with an existing access group, and authentication and privacy keys that a management system can use to access the PacketShaper. This user name is not related to any other user names such as those defined for RADIUS or PolicyCenter access.

Note: If you have not yet defined access groups for your SNMP users, create one or more access groups before you add users to these groups.

Create a New User

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a new SNMP user:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, select **SNMP.**
3. Click the**new**button beside the**Users** heading. The*Add a New User* window opens.**show screen...**

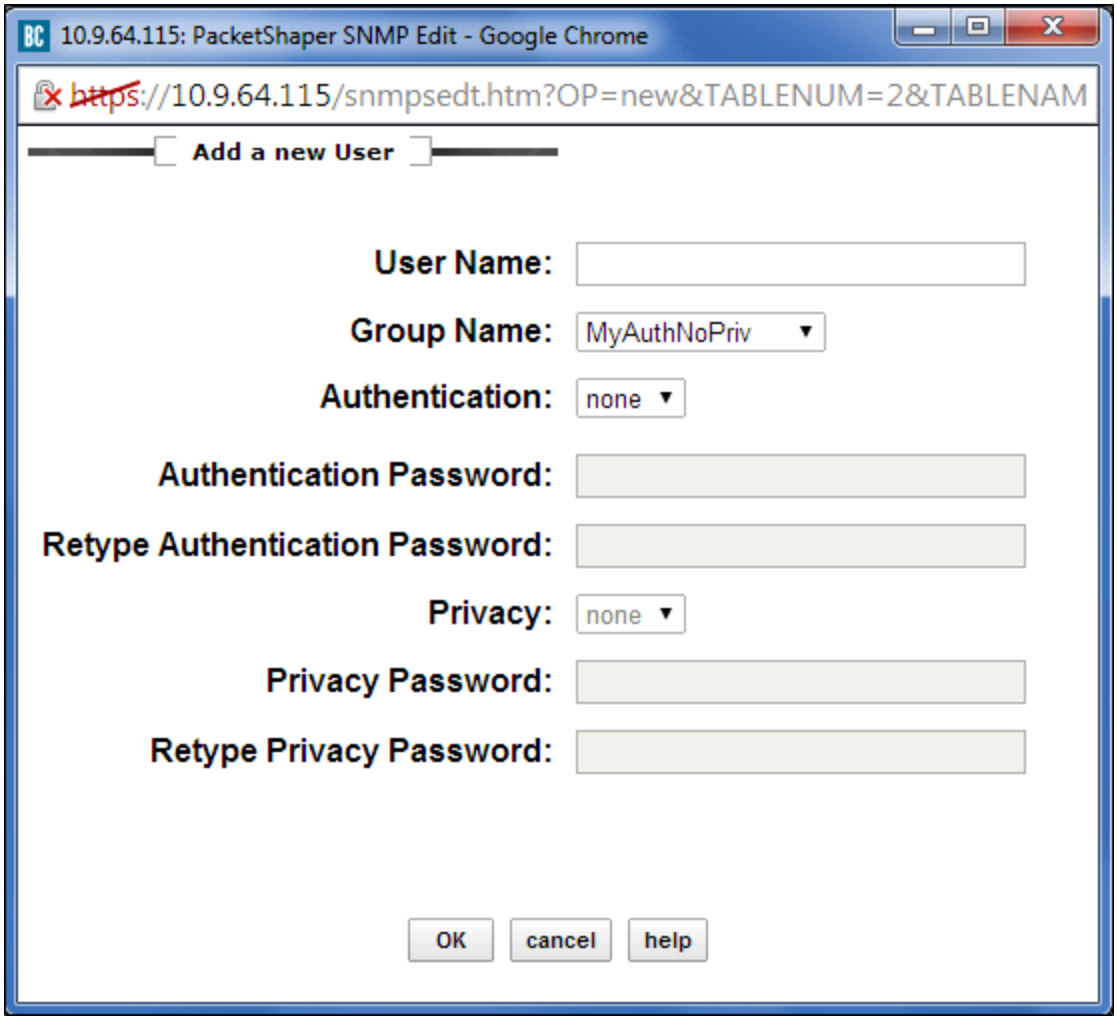

4. Enter a name for the user in the **User Name** field (up to 31 characters; hyphens, underscores, and periods are acceptable).
5. Click the **Group Name** drop-down list and select the access group for the new user.
6. If the new user's access group uses the *usm* (SNMPv3) security model with the *authNoPriv* or *authpriv* security levels, Click the **Authentication** drop-down list and select either the **MD5** or **SHA** authorization protocol. Next, enter an authentication password for the user in the **Authentication Password** and **Retype Authentication Password** fields.

> **Note:** If the user's access group uses the *v1* (SNMPv1) or *v2* (SNMPv2c) security model or the *noAuthNoPriv* security level, keep the default **none** authentication setting.

7. If the new user's access group uses the *usm* (SNMPv3) security model with the *authpriv* security level, Click the **Privacy** drop-down list and select one of the following privacy protection protocols:
   - **des**: CBC-DES Symmetric Encryption Protocol

- **3des**: 3DES-EDE Symmetric Encryption Protocol
- **aes128**: 128- bit AES (Advanced Encryption Standard)
- **aes192**:192- bit AES
- **aes256**: 256-bit AES

  Next, enter an privacy password for the user in the **Privacy Password** and **Retype Privacy Password** fields. If the user's access group uses the *v1* (SNMPv1) or *v2* (SNMPv2c) security model or the *noAuthNoPriv* or *authNoPriv* security levels, keep the default **none** privacy setting.

7. Click **OK** to save your settings. The new user will appear in the *Users* table on the *SNMP* setup screen.

Edit an SNMP User

Modify a user's access group, or change authentication and privacy settings. To edit an SNMP user:

1. If you are not already on the *SNMP* setup page, click the **Setup** tab, and select **SNMP** from the **Choose Setup Page** list.
2. Click the**edit** button beside the*Users* table entry you want to edit. The*Edit User* window opens. **show screen...**

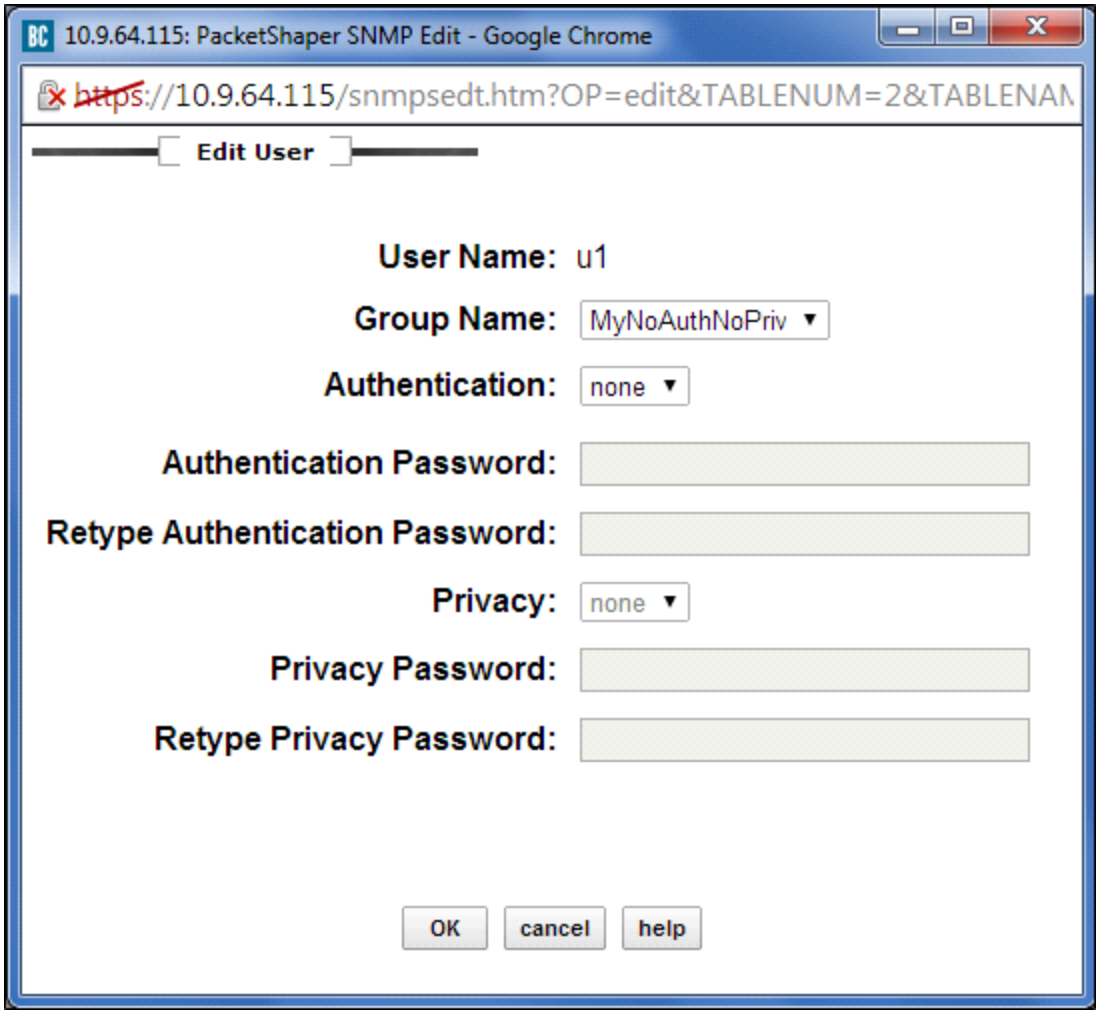

3. Update the user's access group and security settings (described in Create a New User).
4. Click **OK** to save your settings. The updated user will appear in the *Users* table on the *SNMP* setup screen.

Delete an SNMP User

To delete an SNMP user from the *Users* table:

1. If you are not already on the *SNMP* setup page, click the **Setup** tab, and select **SNMP** from the **Choose Setup Page** list.
2. Click the **delete** button beside the user you want to remove. A popup window asks you to verify that the user should be deleted.
3. Click **OK**.

Configure SNMP Notify Tables

To send SNMP v1, v2, or v3 trap notifications to remote users, you define entries in each of the SNMP Notify Configuration tables. Note that each SNMPv3 table supports up to 32 separate entries.

Configure the *Remote Users* Table

A SNMP remote user defines a user or a management system that receives notification of SNMPv3 traps and informs. Unlike a local SNMP user, a remote user is not associated with an access group and therefore has only a notify view, rather than a read or write view.

Create a New Remote User

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a new remote user:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, select **SNMP.**
3. Click the**new**button beside the**Remote Users**heading. The*Add a New Remote User* window opens. **show screen...**

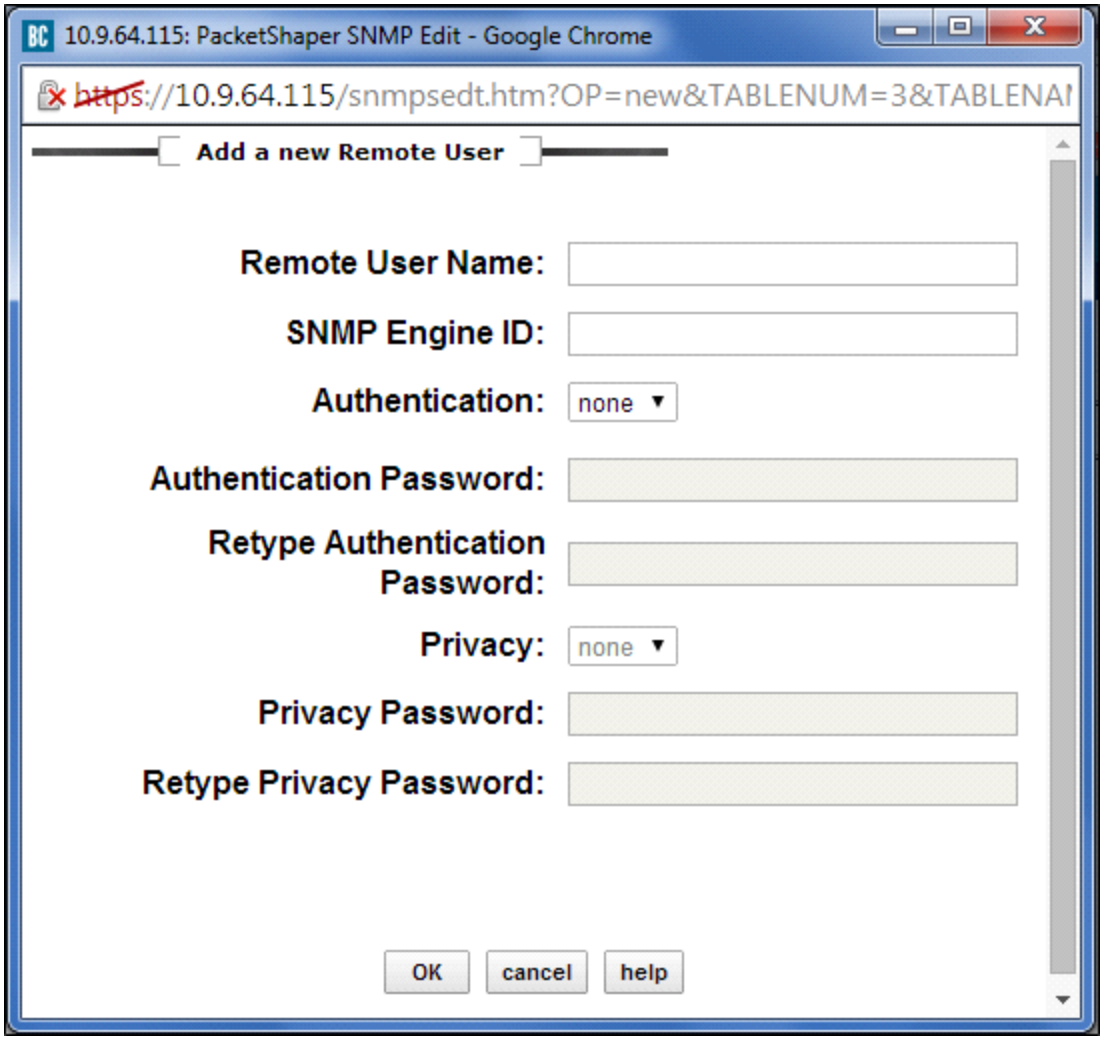

4. In the **Remote User Name** field, enter the name of the new remote user. Remote user names can have up to 32 characters; hyphens, underscores, spaces and periods are acceptable. Each SNMP remote user name must be unique.

5. *(Optional)* An SNMP Engine ID identifies an SNMP engine that will receive trap and inform notifications. The default Engine ID for a remote SNMP user is *LocalSnmpId*, the SNMP agent's own SNMP Engine. If the **SNMP Engine ID** field is left blank, the remote user will user this default *LocalSnmpId* Engine ID. To specify a *different* remote SNMP engine with which this user can communicate, enter the 24-digit hexadecimal Engine ID of a remote SNMP engine in the **SNMP Engine ID** field.

6. If the remote user requires authentication, click the **Authentication** drop-down list specify either the **MD5** or **SHA** authorization protocol. If the remote user does not require authentication, use the default value **none**.

7. If you selected MD5 or SHA authorization in the previous step, enter an authentication password in the **Authentication Password** field and again in the **Retype Authentication Password** field.

8. If the remote user requires privacy protection, click the **Privacy** drop-down list and specify any of the following privacy protection protocols. (If the remote user does not require privacy protection, use the default value **none**.)• **des**: CBC-DES Symmetric Encryption Protocol
   • **3des**: 3DES-EDE Symmetric Encryption Protocol

- **aes128**: 128- bit AES (Advanced Encryption Standard)
- **aes192**:192- bit AES
- **aes256**: 256-bit AES

9. If you selected a privacy protocol in the previous step, enter an privacy password in the **Privacy Password** field and again in the **Retype Privacy Password** field.
10. Click **OK** to save your settings. The new remote user will appear in the *Remote Users* table on the *SNMP* setup screen.

Edit a Remote User

Change authentication and privacy settings, or update a remote user's engine ID. To edit a remote user:

1. If you are not already on the *SNMP* setup page, click the **Setup** tab, and select **SNMP** from the **Choose Setup Page** list.
2. Click the**edit** button beside the*Remote User* table entry you want to edit. The*Edit Remote User* window opens. **show screen...**

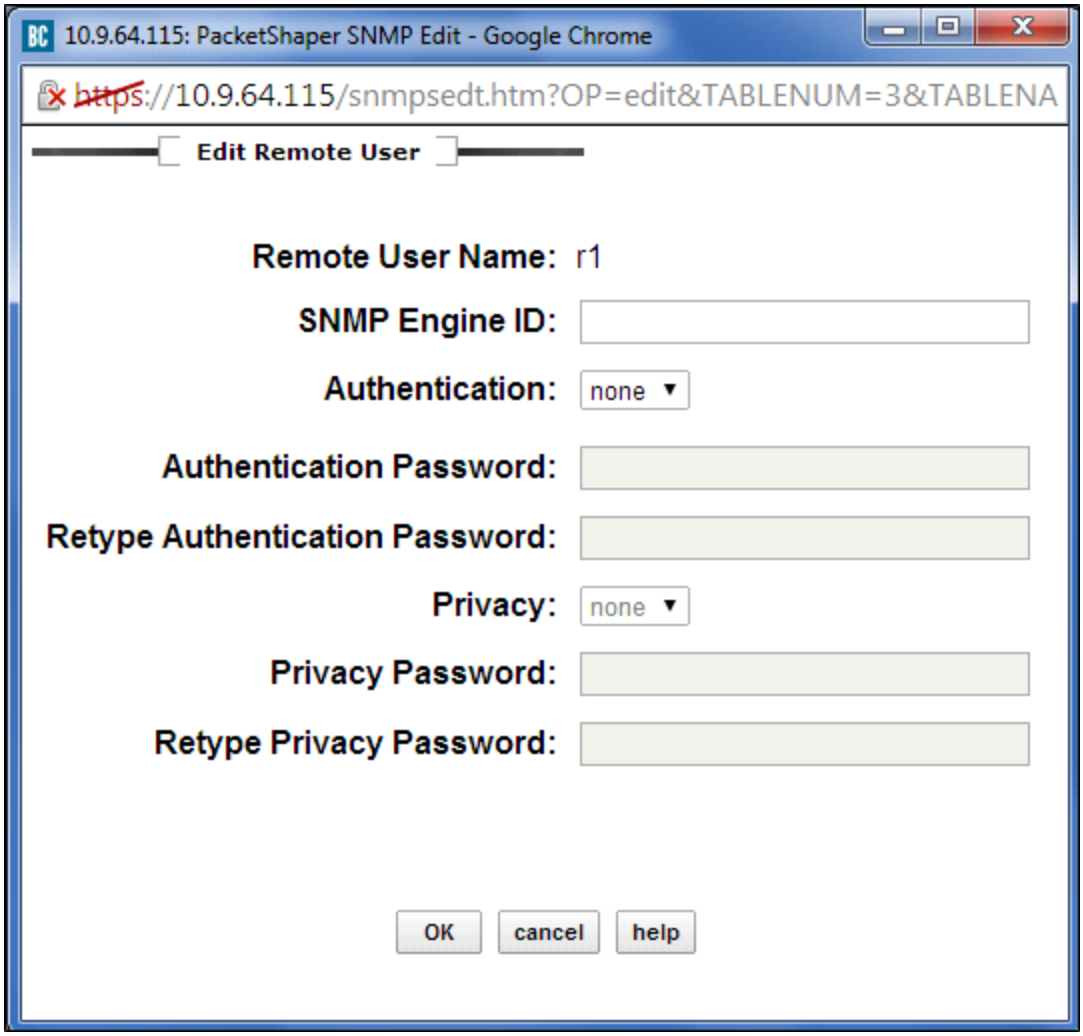

3. Update the remote user's authentication, privacy or engine id settings (described in Create a new Remote User).
4. Click **OK** to save your settings. The updated remote user will appear in the *Remote Users* table on the *SNMP* setup screen.

Delete a Remote User

To delete a remote user from the *Remote Users* table:

1. If you are not already on the *SNMP* setup page, click the **Setup tab**, and select **SNMP** from the **Choose Setup Page** list.
2. Click the **delete** button beside the remote user you want to remove. A popup window asks you to verify that the remote user should be deleted.
3. Click **OK**.

Configure the *Notify Targets* Table

SNMP notify targets determine where SNMP v1, v2, or v3 notifications should be sent (that is, the trap destination). Note that you must define one or more remote users before you can create a target that uses the *SNMPv3* protocol version and the *usm* security model.

> **Note:** If you enter an IP address in the **SNMP Trap Destination** field in the *SNMPv1/v2c Configuration* section, PacketShaper sends traps with the community string *public*. If you want PacketShaper to send traps with a private community string, do not enter the IP address in the **SNMP Trap Destination** field; instead use the **Notify Targets** option in the *SNMPv3 Notify Configuration* section. The **Notify Targets** option can send SNMP v1/v2c/v3 traps with the custom Community Name you configure, as described below.

Create a New Notify Target

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a new notify target:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, select **SNMP.**
3. Click the**new**button beside the**Notify Targets** heading. The*Add a New Notify Target* window opens. **show screen...**

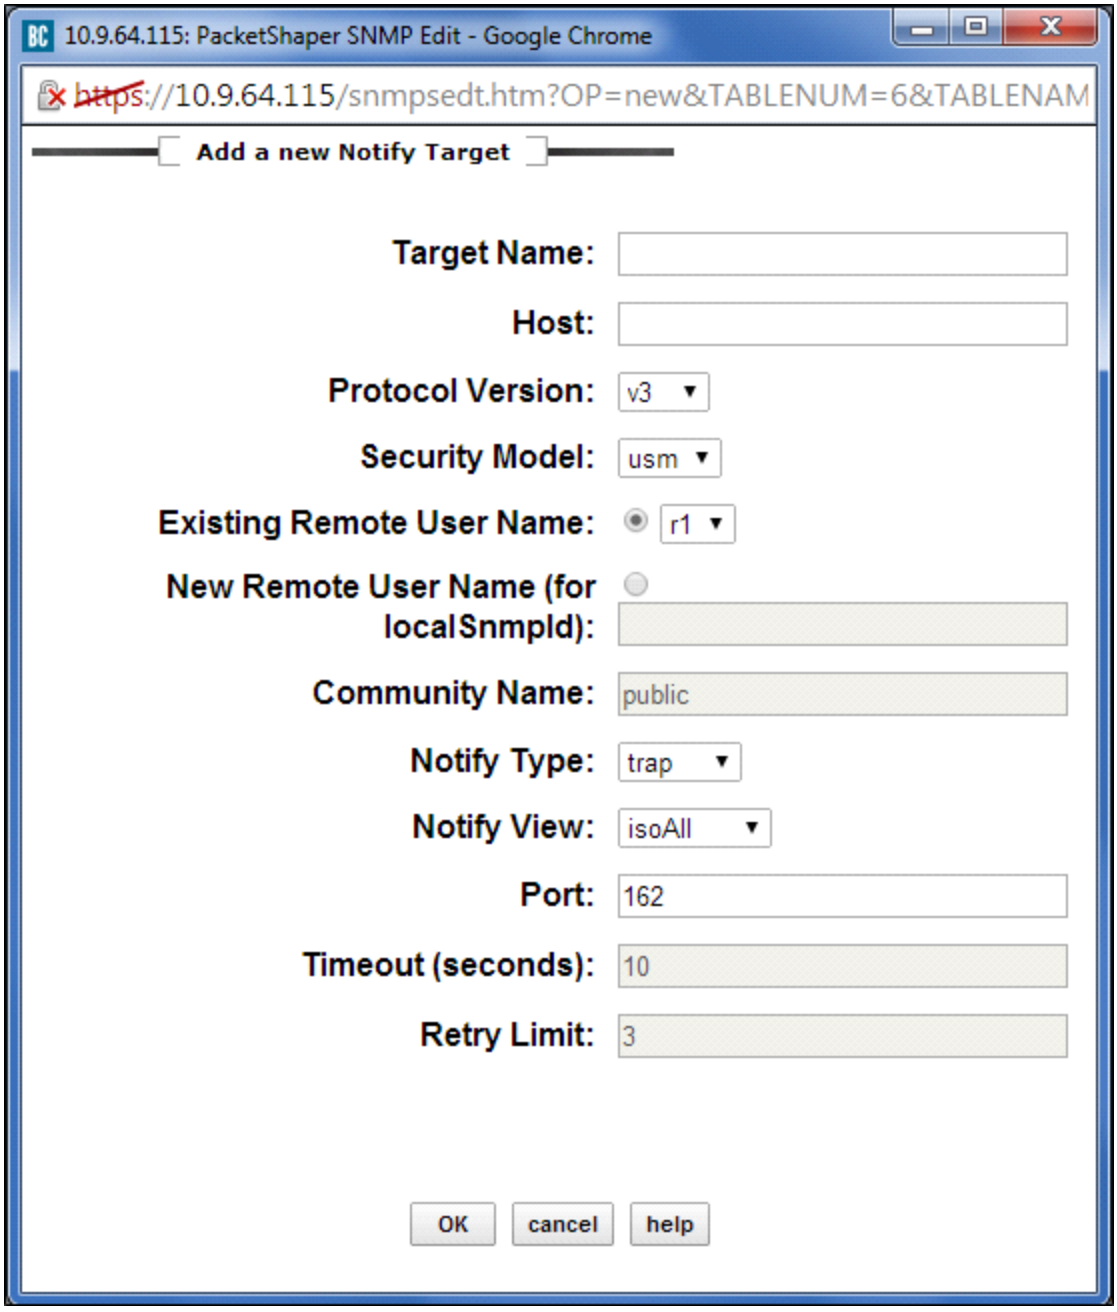

4. In the **Target Name** field, enter a name for the new target. A target name can have up to 32 characters; hyphens, underscores, and periods are acceptable. Each SNMP target name must be unique.
5. In the **Host** field, enter the IP address of a remote IP host (the trap destination), in dotted-decimal format.
6. Click the **Protocol Version** drop-down list and select either **v1** (for SNMPv1), **v2** (for SNMPv2c), or **v3** (for SNMPv3) to specify which version of SNMP notifications the user will receive. If you selected the *v3* option in the **Protocol Version** field, click the **Security Model** drop-down list and specify the security model for this notification by selecting **v1** (for SNMPv1), **v2** (for SNMPv2c), or **usm** (for SNMPv3).

7. If you selected the **v3** protocol version with the **usm** security model, you must also define a remote user for the target.

    - To associate an existing user with this target, select **Existing Remote UserName** and choose an existing remote user name from the drop-down list.
    - To create a new remote user for this target with a localSnmpId and no authorization or privacy protection, enter a name for the new remote user in the **New Remote User Name** field. This localSnmpId user must correspond to a user defined on the remote SNMP machine that will be receiving the traps and notifiations.

8. If your target uses a *v1* or *v2* SNMP protocol or a *v1* or *v2* security model, you can (optionally) enter a new community name in the **Community Name** field. If you don't specify a different community name, the notify target will use the default community name **public**. Allowed characters are: [A-Z] [a-z] [0-9] [_-.,@!$%^*~=+]
9. Click the **Notify Type** drop-down list and select **trap** or **inform** to specify whether the remote user should receive trap notifications or just informs. Note that targets configured to use the *v1* protocol version only support **trap** notifications.
10. Next, select a **Notify View**.
    - To allow the target to receive all types of MIB notifications, use the default Notify View setting **isoAll**.
    - To block the target's access to all data, click the **Notify View** drop-down list and select the predefined view **isoNone**.
    - To limit the remote user's access to a subset of available MIB notifications, click the **Notify View** drop-down list and select a previously defined view.
11. In the **Port** field, enter the port number on the remote host to which the notifications will be sent. The default is port number **162**.
12. In the **Timeout** field, enter the maximum round trip time for communications between the PacketShaper and the SNMP target address, in seconds. Valid timeout values 1-60, and the default value is **10**.
13. In the **Retry** field, enter the number of times the PacketShaper should attempt to retransmit a trap or inform message when it does not receive a response. Valid retry values are 1-10, and the default value is to **3** retries.
14. Click **OK** to save your settings. The new notify target will appear in the *Notify Targets* table on the **SNMP** setup screen.

Edit a Notify Target

Change the parameters of a current notify target, including SNMP protocol and security model settings, remote user or community name. To edit a notify target:

1. If you are not already on the *SNMP* setup page, click the **Setup tab**, and select **SNMP** from the **Choose Setup Page** list.
2. Click the**edit** button beside the*Notify Targets* table entry you want to edit. The*Edit Notify Target* window opens. **show screen...**

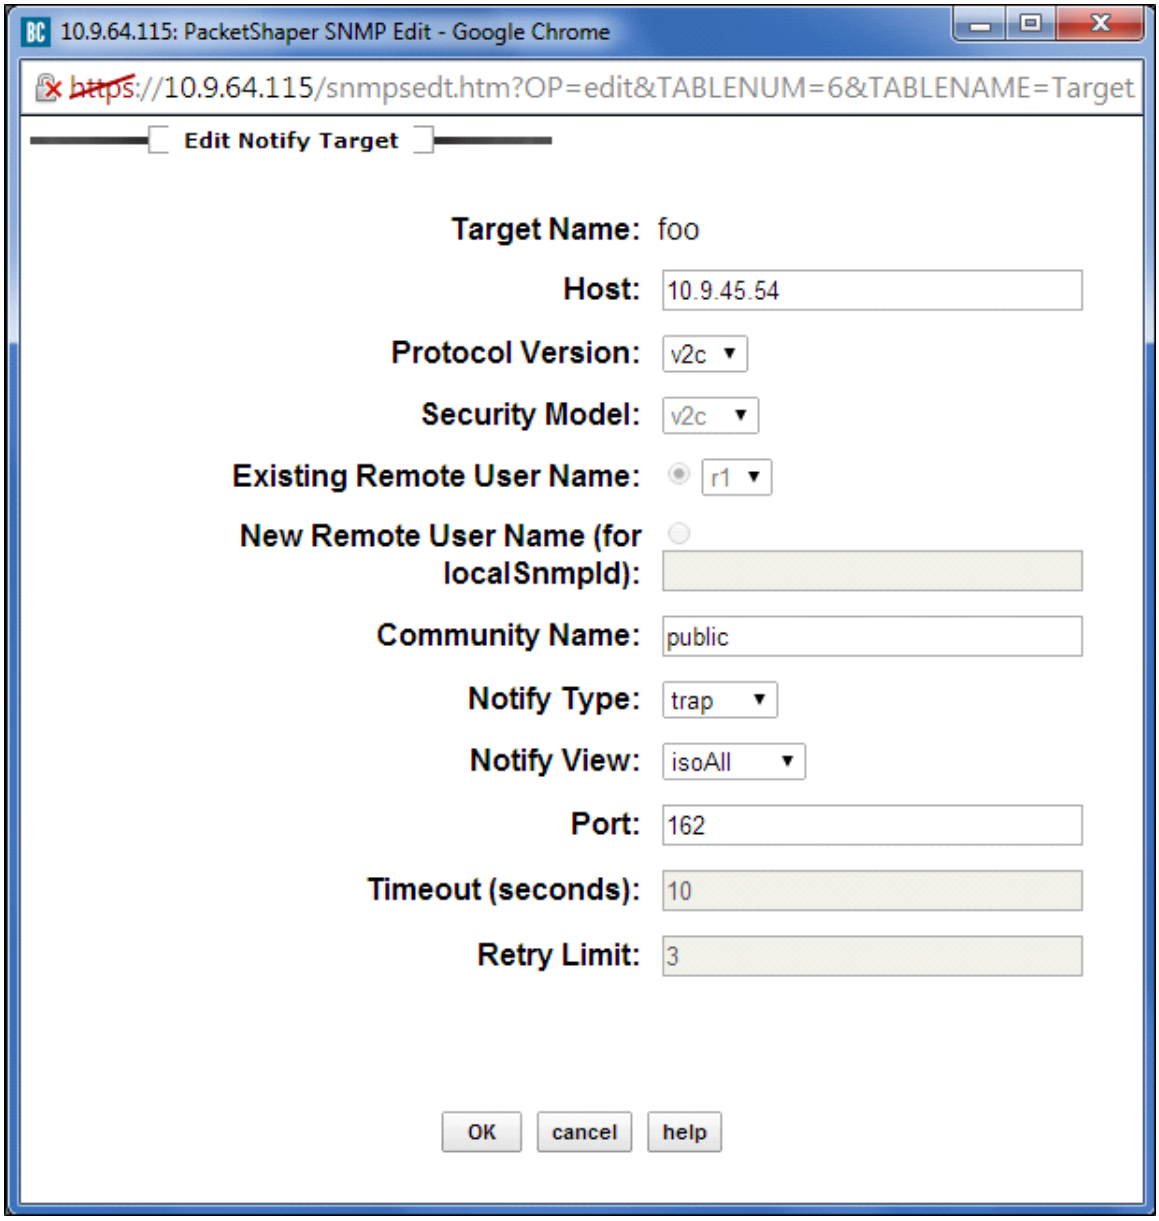

3. Update the values in each field you want to modify (each field is described above in Create a new Notification Target).
4. Click **OK** to save your settings. The updated notify target will appear in the *Notify Targets* table on the **SNMP** setup screen.

Delete a Notify Target

To delete a target from the *Notify Targets* table:

1. If you are not already on the *SNMP* setup page, click the **Setup tab**, and select **SNMP** from the **Choose Setup Page** list.
2. Click the **delete** button beside the notify target you want to remove. A popup window asks you to verify that the target

should be deleted.

3. Click **OK**

Adjust System Variables

PacketShaper's default settings are appropriate for most configurations. However, you can adjust the system variables if your situation warrants it. Use discretion when modifying these variables.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To change a system variable:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **system variables**. The system variable settings appear on the *Setup* screen.
   show screen...

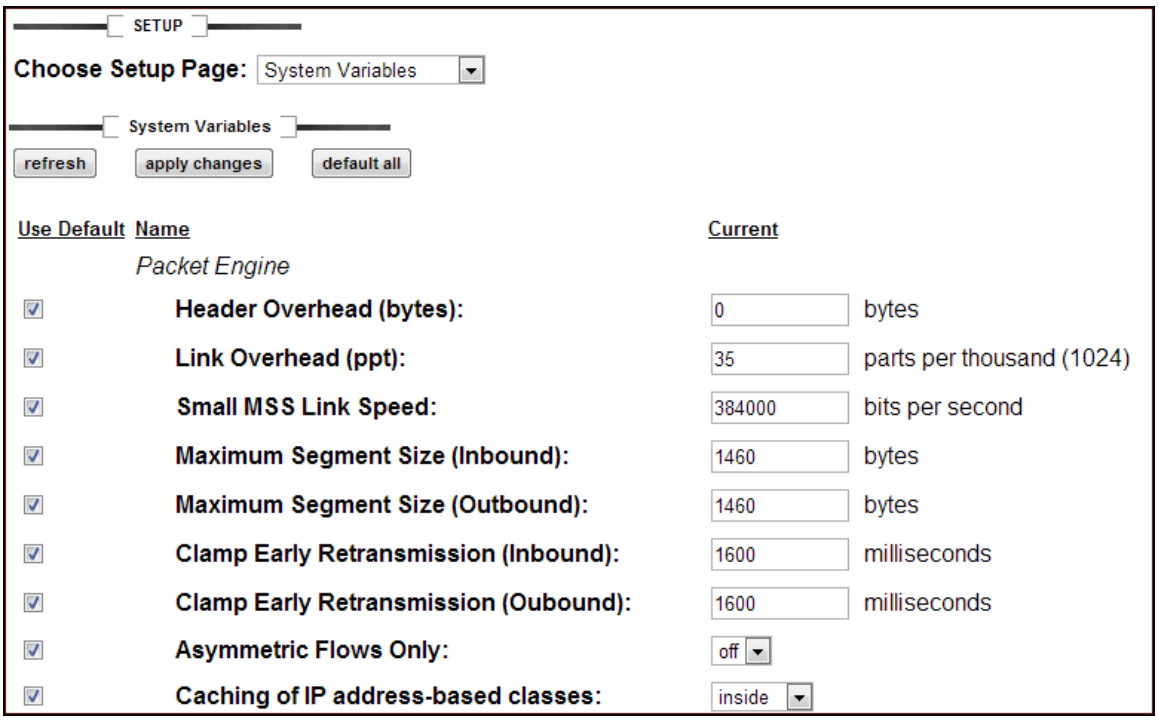

3. Change one or more of the system variable settings (see table below).

> **Tip:** When you place the mouse pointer on a variable name, pertinent information about that variable appears.

4. Click **apply changes**.

5. If necessary, reboot the unit. (Some changes don't take effect until the unit is reset.)

Variable/ Description	Default Value	Min. Value	Max. Value
Packet Engine			
Header Overhead (bytes) Number of bytes that are added to each packet to account for WAN protocol header overhead	0	0	256
Link Overhead (ppt) Number of parts per thousand* by which packet sizes are increased to account for link overhead. This adjustment is useful for links that do bit stuffing. (*Bit stuffing* is the practice of adding bits to a stream of data. Bit stuffing is required by many network and communications protocols, for example to prevent data from being interpreted as control information.) * to be more precise, it s actually parts per 1024	0	0	1024
Small MSS Link Speed Link speeds slower than this value will force the use of smaller MSS (maximum segment size) Prevents PacketShaper from changing the MSS on large WAN links	384000 bps	0	512000
Maximum Segment Size (Inbound) Maximum segment size of TCP packets on Inbound flows. This setting can help avoid packet fragmentation when using VPN and not being able to support 1500-byte packets (the default size) through the VPN tunnel.	1460 bytes	0	65535
Maximum Segment Size (Outbound) Maximum segment size of TCP packets on Outbound flows	1460 bytes	0	65535
Clamp Early Retransmission (Inbound) Number of milliseconds delay for clamping early retransmission time-out on Inbound packets. Puts a maximum on retransmit time.	1600	0 (disable)	3000 (3 sec)
Clamp Early Retransmission (Outbound) Number of milliseconds delay for clamping early retransmission time-out on Outbound packets	1600	0 (disable)	3000 (3 sec)
Asymmetric Flows Only By turning on this setting, PacketShaper will automatically assume all flows are asymmetric and stop TCP Rate Control. In topologies where there are a large percentage of asymmetric flows, this may be more efficient than attempting to apply regular rate control. In addition to disabling rate control, turning on this setting disables all layer 7 classification activities. (PacketShaper must see traffic in both directions in order to classify layer 7.)	off	-	-
Caching of IP address-based classes Cache IP address-based classes on the **inside** or **outside** of the PacketShaper. Change this setting to **outside** to increase performance of classification if the majority of IP addresses in manually created classes are on the outside, rather than the inside.	inside	-	-
Display "503 - Service unavailable" Messages Control the display of the "503 - Service unavailable" server error message when a connection is refused because of admission control (such as a never-admit policy). When set to **off**, the "503 - Service unavailable" message will be customized with the text "This message is sent by Symantec PacketShaper." When set to **on**, PacketShaper will perform a TCP reset and drop the HTTP request; the error message will likely be "The attempt to load http://... failed."	off	-	-

Variable/ Description	Default Value	Min. Value	Max. Value
Enable Support for LFN When enabled, this setting improves performance on Long Fat Networks (LFN) which require larger TCP window sizes. An LFN is a long distance network with large bandwidth and long delay; for example, high-capacity satellite channels are LFNs.	off	-	-
Policy flow limits for all classes Enables/disables the policy flow limit feature. When enabled, PacketShaper will enforce all policy flow limits that have been set on traffic classes. When disabled, all policy flow limits will be ignored. Disabled is the appropriate setting for PacketShapers deployed in proxy or NAT environments. For additional information, see policy flowlimits.	off	-	-
Clip Initial Window Size When this variable is enabled, the PacketShaper will always reduce the initial TCP window size to 1x MSS (maximum segment size). When this variable is disabled, new flows will ramp up faster but enforcement of small rate policies and/or partitions may not work at the beginning of flows.	on	-	-
Auto-discovery			
Non-IP Flows The number of new non-IP connections of a given type that must be identified within a one-minute time frame before PacketShaper creates a class	2	1	1000000
Identifiable Services The number of new connections of an identifiable service to a port less than or equal to 1024 that must be identified within a one-minute time frame before PacketShaper creates a class	1	1	1000000
URL Categories The number of new flows belonging to a particular URL category that must be identified within a one-minute time frame before PacketShaper creates a class for the category	1	1	1000000
Dynamic Ports The number of new connections of an identifiable service to a port greater than 1024 that must be identified within a one-minute time frame before PacketShaper creates a class	2	1	1000000
Static Ports The number of new connections to a static port within a one-minute time frame before PacketShaper creates a Port_#### class in the DiscoveredPorts folder It may be necessary to increase this value on Internet link deployments to prevent excessive number of DiscoveredPorts classes being created. If you don t want any Port_#### classes discovered, set this variable to its maximum value.	100	1	1000000
Create SameSide Class When this variable is enabled, the SameSide class is created automatically. When disabled, the SameSide class will not be auto-created. You may want to disable this variable if traffic is being misclassified into the SameSide class.	on	-	-
Dynamic Partitions			
Active The number of seconds a dynamic partition will be retained after an established flow has sent packets. (Flow is not terminated/completed but hasn't had any recent activity.) **Note:** If no other user needs a dynamic partition, the partition will be retained indefinitely.	300 (5 min)	10	7200 (2 hrs)

Variable/ Description	Default Value	Min. Value	Max. Value
Idle The number of seconds a dynamic partition will be retained after all flows in the partition have been terminated/completed. **Note:** If no other user needs a dynamic partition, the partition will be retained indefinitely.	30	10	7200 (2 hrs)
Reserved for Static The number of partitions reserved for static partitions; all other partitions can be used for dynamic or static partitions	3	0	99
Browser Interface			
Graph Timeout The maximum number of seconds a graph can take to generate in the browser interface; if the graph takes longer to generate than this value, a system time-out error message will appear. **Note:** Increasing this setting can make the browser interface appear to "freeze" while PacketShaper is generating some of the more complex graphs. Sometimes the browser will not display the page until all of the graphs are generated.	60	1	600 (10 min)
Default User Interface The user interface that appears after logging in to the browser interface: the Sky UI or the original (legacy) user interface.	Sky	-	-
Session Management			
Session Purge Timeout This variable controls how many seconds it takes for an unauthenticated login session to get purged from the system. You might need to increase this value if the session times out before PacketShaper can authenticate a login password, for example, when there is latency on the network have not yet been authenticated.	30	30	360 (6 min)
Events			
Registered Events The maximum number of events that can be registered	32	32	128
User-Defined Events The maximum number of events that can be user-defined	32	32	128
Extended SNMP Version Enable/disable the extended SNMP trap for user events. When this variable is turned on, there will be an additional field in the trap that indicates the type of situation that triggered the trap. The field indicates *violated* (when the threshold was exceeded) or *rearm* (when the re-arm value was crossed).	off	-	-
Flow Detail Records (FDR)			
Intermediate FDR Enable/disable the intermediate flow detail records feature. When this variable is enabled, PacketShaper emits intermediate FDRs at the interval specified by the **flowRecordsIntermediateTimeout** variable. **Note:** Enable the intermediate flow detail records feature only when using a suitably-instrumented collector, such as Cisco-based Netflow-5 collectors. Note that IntelligenceCenter ignores intermediate FDRs.	0 (off)	0 (off)	1 (on)
Intermediate FDR Timeout Number of milliseconds between generation and sending of intermediate flow detail records when traffic is present	1500	1000	36000

Variable/ Description	Default Value	Min. Value	Max. Value
Packeteer-P Packets Enable/disable emission of Packeteer-P packets to Packeteer-1 or Packeteer-2 flow detail record collectors. Packeteer-P packets contain statistics that are not related to particular flows, but rather provide information about utilization on the PacketShaper at the time flows are recorded. If this variable is enabled, Packeteer-P records are sent after each UDP flow record packet is sent to Packeteer-1 or Packeteer-2 collectors (not more than once per minute).	0 (off)	0 (off)	1 (on)
Packeteer-P Timeout Number of seconds between generation and sending of Packeteer-0 flow records.	3600	10	5000
Packeteer-0 Packets Enable/disable emission of Packeteer-0 packets to Packeteer-1 or Packeteer-2 flow detail record collectors. Packeteer-0 packets are mapping messages that allow collectors to decipher PacketShaper-related information in the FDRs they receive. For example, in the FDR s ClassID field, a value identifies the traffic class. In order for the collector to understand what class is actually associated with the ID, it uses the class map a list that contains each traffic class on the unit along with the identifying number assigned to each class. If this variable is enabled, Packeteer-0 mappings are sent out approximately once each hour. Note that this variable needs to be enabled only if the collector does not know this information through other means.	0 (off)	0 (off)	1 (on)
Packeteer-0 Timeout Number of seconds between generation and sending of Packeteer-0 flow records.	3600	10	5000
Reset Packeteer 1/2 Counters Controls whether or not the counter fields in FDR packets are reset with each intermediate FDR sent **Note:** This variable only affects Packeteer-1 and Packeteer-2 format FDRs: counter fields are always reset in the NetFlow-5 format.	1 (on)	0 (off)	1 (on)
Voip Metrics			
Enable Latency Calculations Enable/disable the calculation of latency for VoIP metrics. When this variable is enabled, PacketShaper collects data that measure latency for VoIP flows. **Note:** Latency can only be measured between PacketShapers with latency calculations enabled (on).	off	off	on
Latency Probe Send Interval Number of seconds between the issuance of VoIP latency probes that measure VoIP metrics, enabled by the **enableLatency** variable.	5	1	60
Discard Latency Probe Allows the PacketShaper to be configured to discard VoIP latency probes after responding. If VoIP devices located on the Inside of the PacketShaper are sensitive to VoIP latency probes, enabling this variable will prevent potential VoIP call drops.	off	-	-
Users			
Cache Timeout for IPs with User Mapping The number of seconds an IP-user name mapping will be stored in the PacketShaper cache. By default, the user mappings are stored for one hour. Because querying the cache is faster than querying the BCAAA server, you can accelerate user name look ups by increasing the cache timeout. However, the tradeoff is that stale mappings could cause incorrect user name identification.	3600	300	86400

Variable/ Description	Default Value	Min. Value	Max. Value
Cache Timeout for IPs with No User Mapping The number of seconds an IP will be stored in the PacketShaper cache, when the IP lookup does not result in a user.	1800	300	86400
Miscellaneous			
Synthetic Transaction Timeout (Read) Number of seconds after which a synthetic transaction will end when the response received is incomplete	5	1	1000
Synthetic Transaction Timeout (Write) Number of seconds after which a synthetic transaction will be canceled if the server fails to respond to a request	60	10	5000
Link State Mirroring Enable/disable link state mirroring. With link state mirroring, PacketShaper will bring down the second port of a NIC pair if the first goes down. This feature allows each PacketShaper to sit between a WAN router and a switch without blocking detection of switch outages by the router.	off	-	-
Estimate Packet Exchange Time Enable/disable the calculation of packet exchange time. When this variable is disabled, the Pkt Exch column on the Monitor Traffic page will not appear, RTM will not be available, and the packet exchange time and RTM measurement variables will always have a value of 0. After disabling the Estimate Packet Exchange Time variable, you should reset the unit.	on	-	-
Enable Winny Application Classification Enable/disable classification of the Winny service. For optimal performance, enable only when management of Winny traffic is required. **Note:** The Winny peer-to-peer application is used primarily in Japan.	off	-	-
Override Diffserv class sort order Controls the sort order of the traffic tree, with respect to Diffserv classes (those with DSCP marks). Three settings are available: **0** Diffserv classes are sorted below IP-address-based classes, but above port-based classes (the default). **1** Diffserv classes are sorted above IP-address-based classes. **2** Legacy sort order (Diffserv classes are sorted after IP-address-based classes, port-based classes, and auto-discovered classes.) **Note:** The new sort order doesn't take effect until the unit is rebooted.	0	0	2

Revert to Default System Variable Settings

You can revert a single system variable to its default setting or revert all system variables to their defaults.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

If you revert a PolicyCenter child configuration to its default system variable settings, the child configuration will inherit all its system variable settings from its parent configuration. If you issue this command from a root-level configuration, that configuration will reset its values back to the default settings.

To change a system variable:

1. Click the **Setup** tab.

2. Display the *System Variables* setup screen. **show screen...**

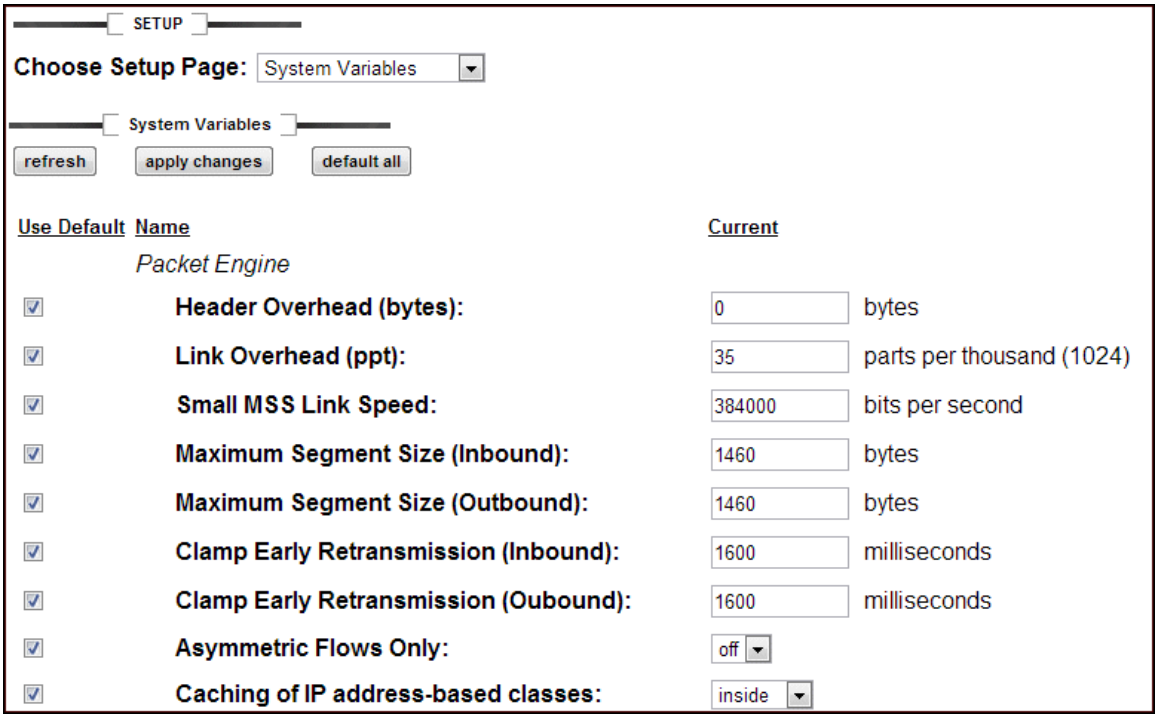

3. To revert one variable to its default, click the **Use Default** checkbox next to the variable name.

or

To revert all variables to their default settings, click **default all**.

4. Click **apply changes** to update the settings.

Configure Heartbeat Emission

PacketShaper emits a message, called a *heartbeat*, to a designated Symantec server on the following occasions:

- PacketShaper bootup (on warm reset or cold boot)
- Daily (every 24 hours of uptime)
- After a system failure (on bootup, following a system restart)

The heartbeat message provides Symantec support professionals with key information about the PacketShaper, such as what modules are enabled, its memory allocation and usage, how long the unit has been running since last reboot (uptime), traffic load conditions, and basic configuration settings. Bootup and crash heartbeat messages append the boot log (a list of the versions that were booted on the PacketShaper, including the date and time of each boot). Crash heartbeat messages include the crash log.

Using the information contained in the heartbeat messages, Symantec is able to provide better, faster support to its users. It can also compile statistics on the stability of various software releases and hardware products. The heartbeats can also be used to identify and resolve defects.

> **Note:** Although you can configure heartbeats in PolicyCenter, PolicyCenter itself does not send the heartbeats. Only PacketShaper devices emit heartbeats.

Heartbeat emission is enabled by default. Symantec recommends that you *not* disable the feature. Be assured that the messages are encrypted, contain no private information such as passwords and IP addresses, and are sent securely via HTTPS. The size of the daily heartbeat message is negligible (30-40K) and has virtually no impact on PacketShaper performance.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To configure heartbeat emission:

1. Click the **Setup** tab.
2. From the setup page list, select **heartbeat**. The heartbeat settings appear on the *Setup* screen. **show screen...**

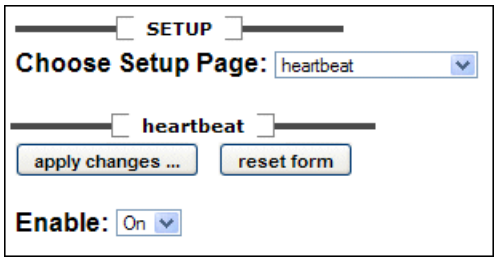

3. Select **On** for **Enable** to enable heartbeat emission or **Off** to disable.
4. Click **apply changes**.

Note: If the Symantec server is unreachable due to network or DNS resolution issues, the heartbeat message will *not* be sent. There is not a retry mechanism. PacketShaper will attempt to emit the heartbeat at the next 24-hour interval.

Configure URL Categorization

To configure URL categorization, you can:

Enable WebPulse

Select Categories to be Discovered

Download the Latest WebPulse Maps

Manually Update the WebPulse Maps

PacketShaper uses mapping files to integrate with WebPulse, to identify URL category names, web application names, and operation names. The WebPulse maps support the following features:

- **URL categorization:** The category map associates URL category names with numeric IDs. PacketShaper refers to the map to look up the category name after WebPulse has assigned a category ID to a flow.
- **Web application identification:** The web application map associates web application names with numeric IDs, and lists the operation IDs applicable to the application. PacketShaper refers to the map to look up the application name after WebPulse has assigned an application ID to a flow. In addition, PacketShaper uses the map to determine which operations are applicable to an application, for example, to create a class for a specific operation for a web application (such as Facebook-Post_Messages, Gmail-Upload_Attachment).

When you enable WebPulse, the latest maps are automatically downloaded to your PacketShaper. Thereafter, PacketShaper automatically downloads the map files every day and after a device reset, regardless of whether the files have changed or not. If you want to make sure your PacketShaper has the latest WebPulse maps, you can use the **update now** button on the WebPulse setup page.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To manually update the WebPulse maps:

1. Click the **Setup** tab.

2. Choose the **WebPulse** setup page.

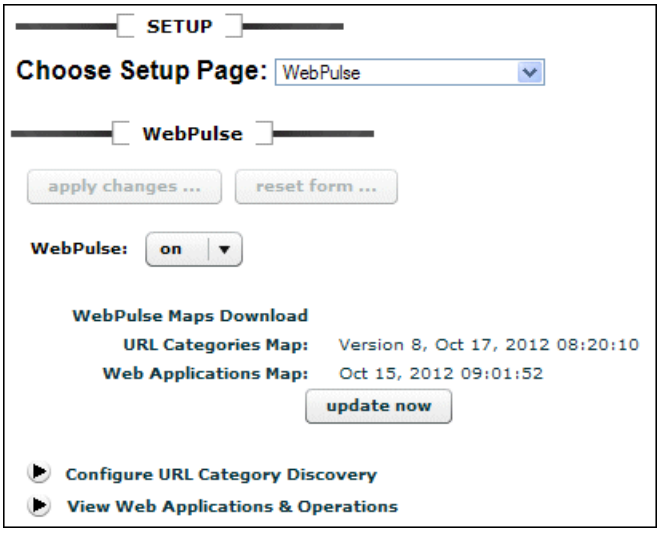

The version of the WebPulse maps are displayed, as well as the date of the last download. For example:

URL Categories Map: Version 8, Oct 10, 2014 09:49:10
Web Application Map: Oct 10, 2014 009:49:10

3. Click **update now**.

The map files are downloaded, replacing the current files (even if the files haven't changed). The version number will change if applicable and the time stamp will reflect the date and time the file was updated.

> **Note:** The Categories map version is incremented when a category has been added, renamed, or deleted.

Select Categories to be Discovered

If you only care about monitoring and controlling some of the URL categories, you can turn off the auto-discovery of the categories you are not interested in. By doing so, you will not clutter your traffic tree with classes you don't care about. By default, discovery is enabled for all categories.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To select categories to be auto-discovered:

1. Click the **Setup** tab.

2. Choose the **WebPulse** setup page.

3. Click the next to **Configure URL Category Discovery**. The categories are listed.

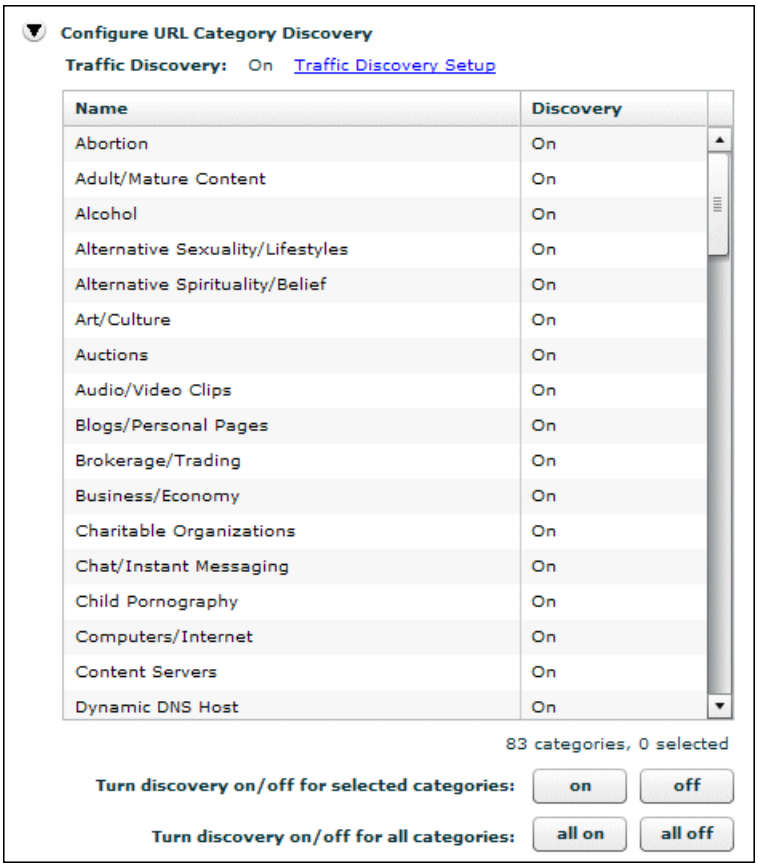

4. In the *Configure URL Category Discovery* table, hold down Ctrl and click each category you *do not* want to discover; then click **off.**

or

If you want to discover only a few categories, click **all off** to set discovery off for all categories. Then select the categories you *do* want to discover and click **on**.

5. Click **apply changes**.

Additional Information

- Even when you have turned off discovery for a category, you can still create a class based on this category and PacketShaper will classify traffic into the category.

- By default, PacketShaper auto-discovers category classes after it sees one flow for the category in a one-minute time interval. This threshold is user configurable via the **URL Categories** system variable. See Adjust System Variables in PacketGuide.

Configure WebPulse

WebPulse is a service provided by Symantec that receives URL identification requests and responds with the IDs of the category, application, and/or operation associated with the URL. When WebPulse is enabled, you will have the ability to:

- create classes based on specific URL categories, and PacketShaper will then classify web traffic that corresponds to each of these categories into the appropriate class
- auto-discover category classes (you must first create an "all categories" class and enable class discovery)
- create classes based on specific web applications and operations
- auto-discover web application classes

The WebPulse URL database is either in the cloud or stored locally on the appliance. Note that the option for a locally stored database is available only on PS-S500 appliances using WebPulse Express (see "setup webpulse express" on page 1367).

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To configure WebPulse:

1. Click the **Setup** tab.
2. Choose the **WebPulse** setup page.

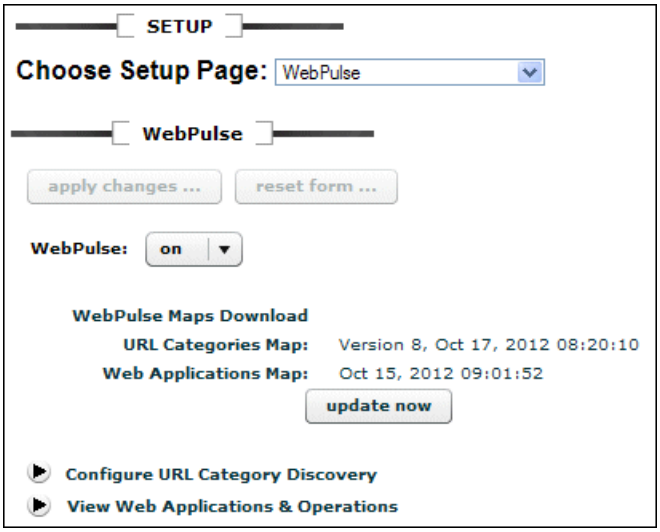

3. For **WebPulse**, select **on**.
4. *(Optional)* . Select which categories you want to auto-discover.
5. Click **apply changes**.

Additional Information

- By default, PacketShaper auto-discovers category classes after it sees one flow for the category in a one-minute time interval. This threshold is user configurable via the **URL Categories** system variable.

- If you have secured the PacketShaper's outside interface, the WebPulse feature requires that you add WebPulse servers to the exception list. In addition to the servers listed on "PacketShaper Port and URL Usage " on page 924, be sure to add the fastest WebPulse service points (use the **setup webpulse show service** CLI command to find the IP addresses).

- Although you can create URL category and web application classes when WebPulse is disabled, there won't be any classification into these classes until you enable WebPulse.

- If WebPulse is disabled, PacketShaper does not look up URLs in the WebPulse cloud, in the URL cache, or in the WebPulse Express local database.

Configure User Awareness on PacketShaper

Before you can create user-based classes or monitor user activity on your network, you need to perform the following one-time configuration tasks:

- Install BCAAA
- Configure PacketShaper as a BCAAA Client

Install BCAAA

PacketShaper uses the Blue Coat Authentication and Authorization Agent (BCAAA) to resolve IP addresses to user names so that it can classify and report on network users and groups. You must install BCAAA version 6.1 on a Windows Domain Controller (DC) or other member server (running version Windows 2008 or 2012). When deciding where to install BCAAA, keep in mind that it must be able to contact every DC that contains user login sessions for the users you are authenticating. Therefore, you should try to keep BCAAA on the same subnet as the DC whenever possible to avoid excess WAN bandwidth usage.

Additional Information

- You must have a MySymantec account with a valid support contract in order to download BCAAA.
- BCAAA version 6.1 is required.
- The BCAAA-PacketShaper communication port on your Windows server firewall needs to be opened up for BCAAA to work. The BCAAA Setup wizard has an option for configuring this setting.

Planning for your BCAAA Installation

It's important to size your BCAAA installation appropriately when preparing to use the user awareness feature. To ensure that BCAAA can service all the queries from PacketShaper, Symantec recommends the following BCAAA-to-PacketShaper ratios:

- One instance of BCAAA per PacketShaper for any number of supported hosts.

  or

- For redundancy, a PacketShaper can connect to two BCAAA servers. If PacketShaper detects that the primary BCAAA server is down, it automatically connects to the secondary BCAAA server. One server is active, and the other is standby.

or

- Multiple PacketShapers can share a single BCAAA server if the number of active hosts on each PacketShaper is less than or equal to the IP-user cache size for that model (equivalent to three times the supported number of users).

Disregarding these recommendations could overload the BCAAA server, resulting in BCAAA being unable to respond to PacketShaper IP-user queries. See Configuration Limits in PacketGuide for per-model specifications.

> **Caution:** To prevent any impact on ProxySG authentication functionality, PacketShaper should not share a BCAAA server with a ProxySG.

Download and Install BCAAA

The user who installs BCAAA needs admin rights on the BCAAA server. Installation can be performed by a local admin – that is, an account that exists only on the BCAAA server, and not in Active Directory. The install could also be performed by an Active Directory admin.

1. Log in to the Windows server where you plan to install BCAAA.
2. Go to the Symantec Support site, MySymantec.com, and log in using your MySymantec credentials.
3. In the products list, locate the serial number of the appliance.
4. Download the BCAAA Windows zip file for the desired version.
5. Unzip the BCAAA setup package and double-click the .exe file to launch the BCAAA Setup wizard.
6. To begin the setup, click **Next**.
7. **Destination Folder:** Specify the location for the BCAAA software. You can accept the default location (C:\Program Files\Blue Coat Systems\BCAAA) or browse to a different location. Make sure that your anti-virus software is not configured to scan the directory in which you install BCAAA. Click **Next** to continue.

   > **Note:** If you are installing on a system that is running a previous BCAAA version, make sure you install to the same location as the previous version to ensure that your configuration settings are retained.

8. **Port Number:** Specify the port number that BCAAA and the PacketShaper will use to communicate. By default, both BCAAA and the PacketShaper use port 16101. If you choose a port other than the default, you must set the same value on the PacketShaper. In addition, make sure that this port is not blocked by a firewall between the BCAAA server and the PacketShaper or by the Windows firewall on the server where you are installing BCAAA. Click **Next** to continue.
9. **Firewall Configuration:** Enable the **Open port** option to have the BCAAA Setup wizard open the BCAAA-PacketShaper communication port (16101 by default) on the Windows server firewall. If you don't enable this option, you will need to manually configure this setting yourself.
10. **SSL Requirements:** When prompted whether you want to use SSL between the PacketShaper and BCAAA, select **Forbidden**, and then click **Next**

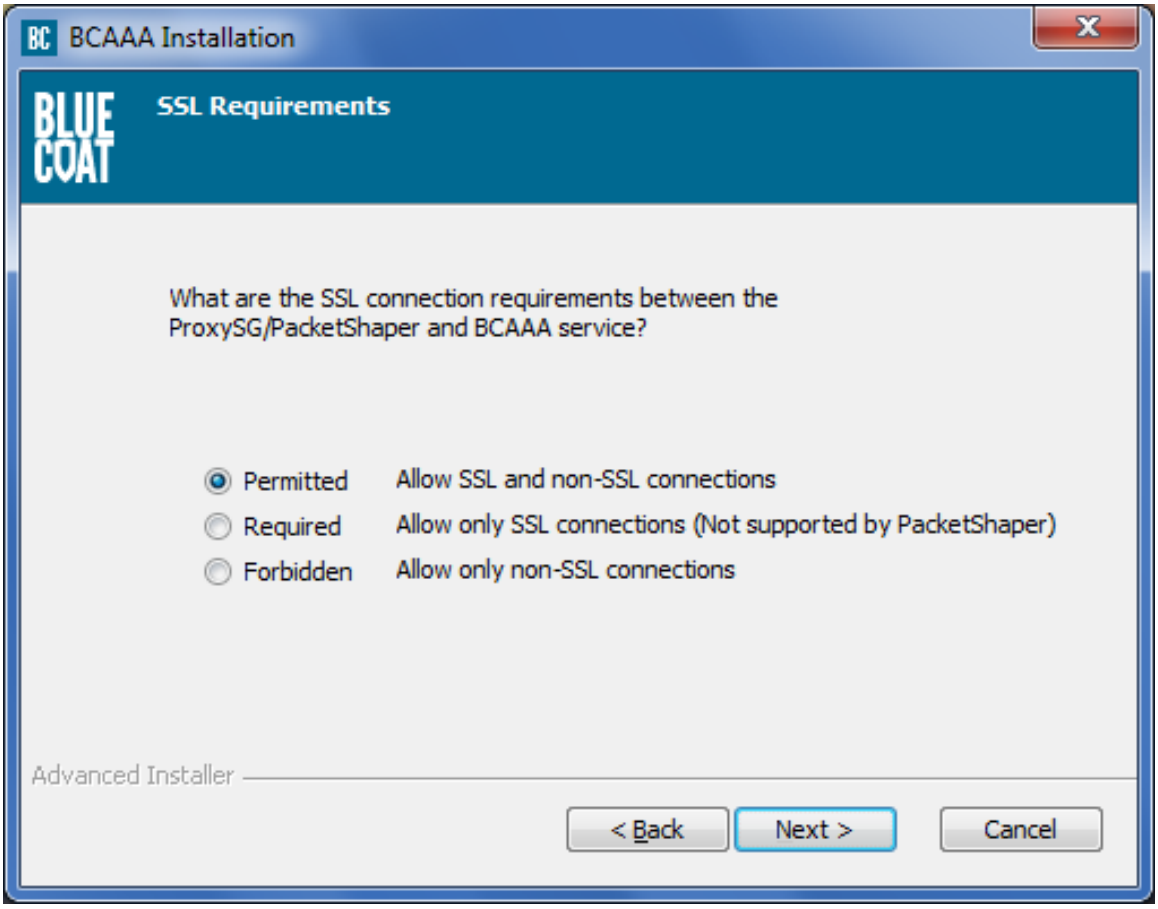

11. **Windows SSO Configuration:** When asked whether this installation will be supporting a ProxySG with a Windows SSO realm or a PacketShaper, select **Yes**. Click **Next** to continue.
12. **Service Account Configuration:** Enter the user name and password that the BCAAA service will log on as. The account must be a member of the domain. Click **Next** to continue.
13. Click **Install**. When installation completes, the final BCAAA screen displays.

> **Note:** To uninstall the BCAAA agent, double-click the BCAAA .exe file to launch the BCAAA Setup wizard, then select **Remove**. Note that if you need to modify a setting after installing BCAAA, you must remove it and then re-install.

Configure PacketShaper as a BCAAA Client

PacketShaper uses the Symantec Authentication and Authorization Agent (BCAAA) to resolve IP addresses to user names so that it can classify and report on network users and groups. The first step is to install BCAAA on a Windows Domain Controller or other member server. Once BCAAA is installed, you can configure PacketShaper as a BCAAA client, as described below.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To configure PacketShaper as a BCAAA client:

1. Click the **Setup** tab.

2. From the list of setup pages, select **User Awareness**. The User Awareness setup screen appears.

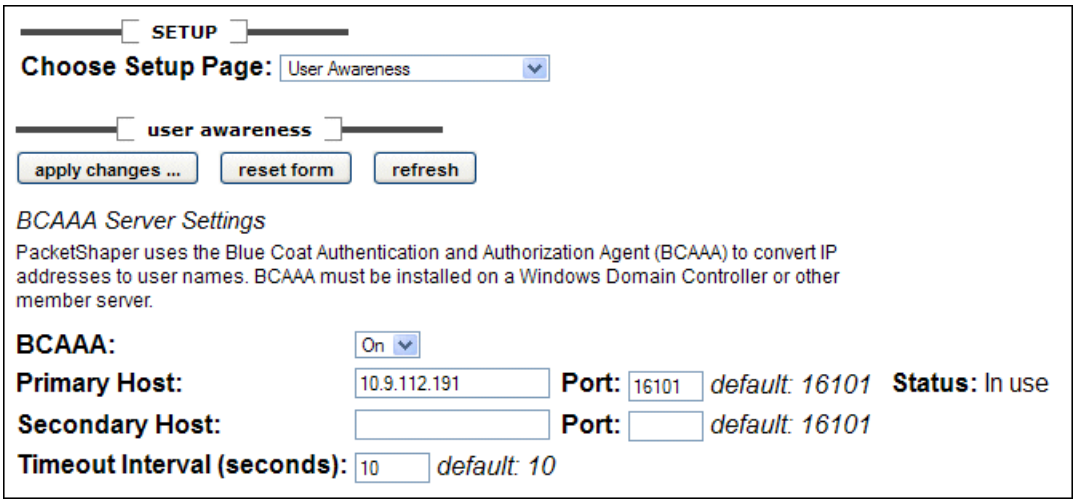

3. For **BCAAA**, select **on** to enable communication with BCAAA.

4. For the **Primary BCAAA Server**, specify the following details:

   - **Host Name** – IP address or domain name of the primary BCAAA server
   - **Port** – Port number. If not specified, the default port is used (16101). This port number must match the port number configured on BCAAA.

5. (Optional) Configure a **Secondary BCAAA Server**. PacketShaper will automatically use the secondary server if the primary server is unreachable.

6. For **Timeout**, enter the number of seconds for PacketShaper to wait for a response from the BCAAA servers when looking up a user name (2-30 seconds, default=10).

7. Click **apply changes** to save the settings.

8. Verify PacketShaper can connect to the BCAAA servers.

Verify BCAAA Server Configuration

After you have configured the BCAAA server(s) and applied the changes, you can verify the server configuration by checking the status on the User Awareness setup page.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To check the status of BCAAA servers:

1. Click the **Setup** tab.

2. From the list of setup pages, select **User Awareness**. The status appears on the same line as the server settings.

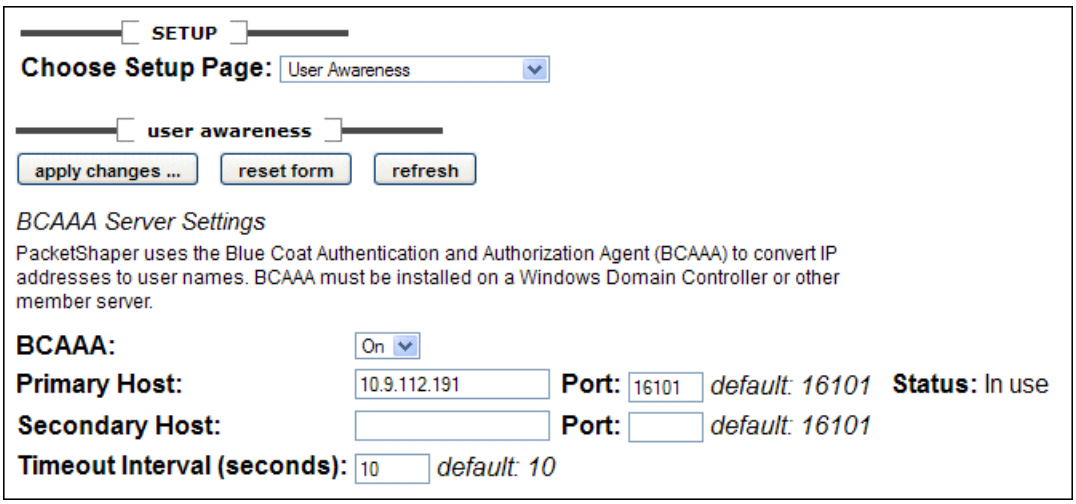

3. Click **refresh** to update the status; note that it may take several minutes after configuring a server before the final status is displayed.

Status	Description
Attempting to connect	PacketShaper is in the process of trying to establish a connection to this BCAAA server.
In use	PacketShaper is connected to this BCAAA server and is using it to look up users.
Available	PacketShaper is able to connect to this server but is using another working server (one with an "In use" status) for user lookups.
Failed connection	PacketShaper tried to connect to this BCAAA server but the connection failed. Possible reasons for a failed connection include: ■ An incorrect IP address is specified. ■ BCAAA is configured to use a different port than PacketShaper. ■ BCAAA server is down. ■ Link is down. ■ No route exists between the networks.

Once you have a BCAAA server that shows an "In use" status, you can then create user-based classes, create user group classes, view pie charts of top ten users, list top users of a certain service, and perform host analysis on a per-user basis.

Status messages are also displayed on the Info page in both the Legacy and Sky UI. For example:

User Awareness: Primary BCAAA Server, host 10.9.112.195, connection failed.

Or, if BCAAA is enabled but a server hasn't been configured, the message is:

User Awareness: No BCAAA server configured.

Another way to verify the server configuration is by testing whether an IP address returns the correct user name listed in Active Directory; use the**setup bcaaa test** CLI command.

Traffic Classification Tasks

Before you can manage bandwidth allocation, you must determine exactly what applications are running on your network. PacketShaper s traffic discovery feature can automatically provide you with a list, called a *traffic tree*, of all the applications on a network. PacketShaper differentiates one application from another by evaluating characteristics in traffic flows and grouping them into *classes*. Each traffic class contains at least one *matching rule*, which is a set of characteristics that identifies a specific traffic type. For more information, see the Traffic Classification Overview.

PacketShaper can automatically create classes as traffic is discovered on your network, or you can create classes manually. This section includes tasks for enabling automatic traffic discovery and creating traffic classes manually.

Enable/Disable Traffic Discovery

When traffic discovery is enabled, PacketShaper automatically creates traffic classes based on the service types it detects. When traffic discovery is disabled, traffic classes are not autodiscovered.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To turn traffic discovery on or off:

1. Click the **Setup** tab. The *Basic Settings* screen is automatically displayed in the *Setup* window.

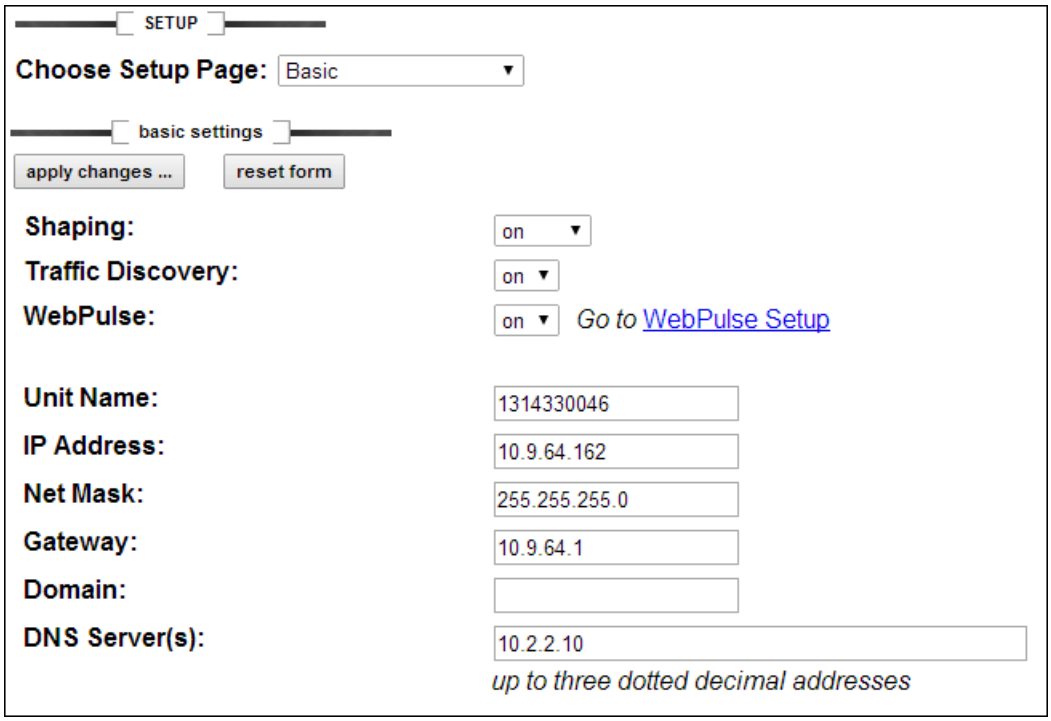

2. In the **Traffic Discovery** field, select **on** to enable discovery or **off** to disable.

3. Click **apply changes** to update the settings.

Turn on Traffic Discovery for an Individual Class

Certain classes such as PVC, VLAN, subnet, Citrix, DICOM, Oracle, PostgreSQL, RTCP, and RTP classes can be more specifically classified. For these types of classes, you can choose to enable traffic discovery within the class. If you turn on traffic discovery:

- for a subnet class (such as Marketing or Engineering), PacketShaper will discover all the traffic running on the subnet and create child classes for these services.
- for Oracle and PostgreSQL classes, PacketShaper will create child classes for each database that is discovered.
- for Citrix classes, PacketShaper will create child classes by published application.
- for DICOM classes, PacketShaper will automatically discover DICOM client or server applications.
- for RTP and RTCP classes, PacketShaper will automatically create child classes by encoding name (codec). Or, if you rather sub-classify RTP traffic by user agent, you can enable the enableVoIPUseragentAutoDiscovery system variable.
- for ProxySG-ADN classes, PacketShaper will automatically create child classes for applications or ports used in the ADN tunnel.

Traffic discovery cannot be enabled within some classes because the auto-discovery process cannot categorize those traffic types in any greater detail. For example, a class that matches on IPX cannot be further auto-classified.

In order for class discovery to work, you must have traffic discovery enabled. If discovery is enabled for a class, but is turned off for the unit, new classes will not be discovered. See Enable/Disable Traffic Discovery.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To turn on traffic discovery for a specific traffic class:

1. Click the **Manage** tab on the navigation bar. The *Traffic Class* window appears.

2. Select the class name In the left window pane. The class' details appear in the *Traffic Class* window.

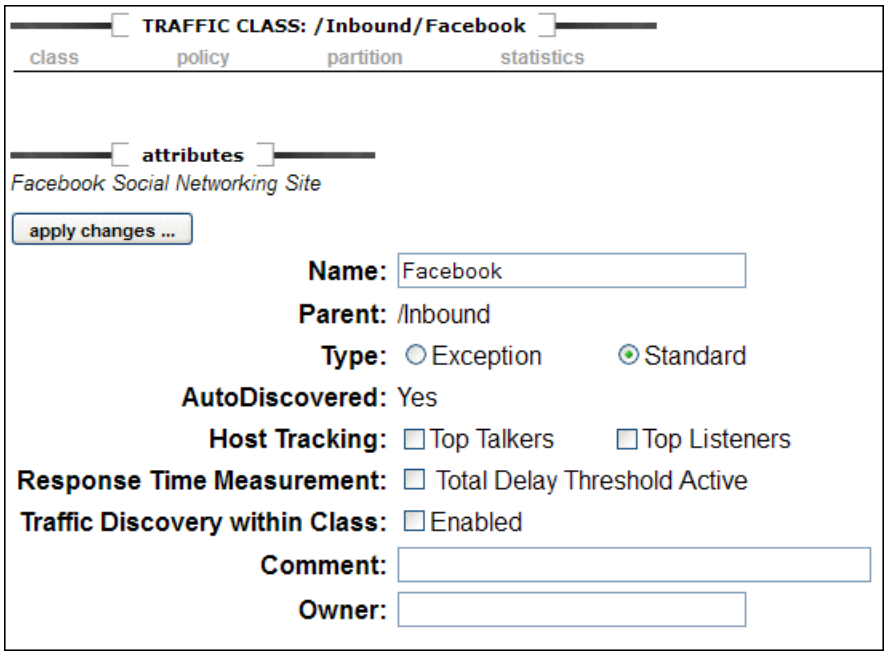

3. Locate the **Traffic Discovery within Class** option.

> **Note:** If "Not Available" appears next to **Traffic Discovery within Class**, the selected class cannot be classified in any greater detail.

4. Select the **Enabled** checkbox.

5. Click **apply changes**.

Create a Traffic Class

While traffic discovery creates traffic classes for the applications and services it identifies, there may be times when you will want to create a class to handle a more specific traffic type. For example, you may need to create classes for subscribers or branch offices. Or perhaps you might want to monitor and control the traffic for a particular website. In this case, you can manually create a child class under the HTTP parent class and specify a domain.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To define a new traffic class:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder that will be the parent of the class you are about to define.

3. Click **class** and select **add**. The *New Traffic Class* screen appears. **show screen...**

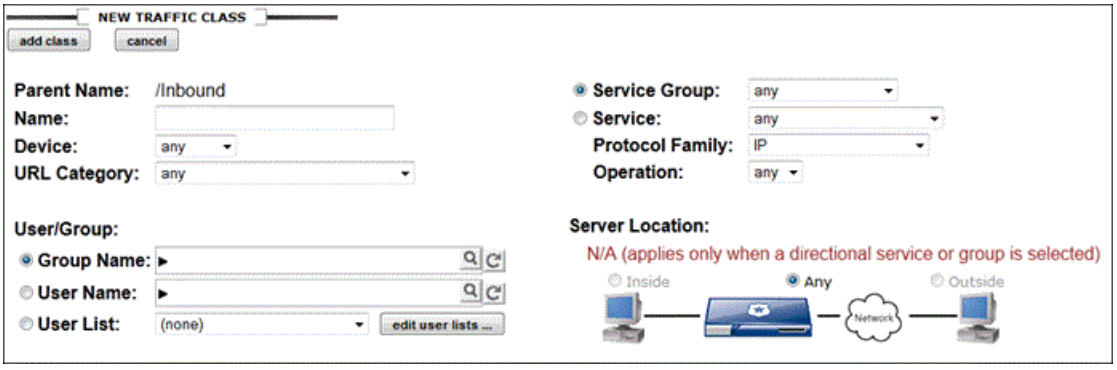

4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).

5. *(Optional)* To classify by a particular pair of interfaces on a slot (such as Slot2), select the Device.

> **Note:** Leave **Device** set to **any** unless you want to classify a LAN separately from the rest of the network traffic.

6. To classify by a particular URL category, choose the URL Category from the drop-down list. (See Create a Category-Based Class and Create a Class for Multiple Categories.)

If you don't want to classify by URL category, keep the default setting (**any**). To create a class that discovers child classes for URL categories, choose **All Categories**, and then turn on class discovery.

7. Choose a **Protocol Family**. For most classes, the default protocol family, **IP**, is correct. To create a class for IPv6 traffic, however, you must select **IPv6**; once you do this, the **Service** list (step 8) then displays the IPv6-based services.

8. Choose a **Service** or **Service Group**. A *service* is a particular application or protocol that PacketShaper can identify and classify. A *service group* is a set of services that have been grouped together based on a common functionality. For a complete

list of services and service groups built in to PacketShaper, see Applications, Protocols, and Services.

Child classes must have the service explicitly set to the service of the parent rather than relying on the service defined in the parent class.

9. Complete the entries in the *New Traffic Class* window to define the matching rule for the class. For descriptions of these fields, see Matching Rules.

10. Click **add class**.

When you add a class to a parent that has no children, a Default, match-all class is created as the sibling of the newly added class. If there is an existing policy for the parent, it is transferred to the Default class since a parent class cannot have a policy. If you later delete all the child classes, the policy will be transferred automatically back to the parent class.

> **Note:** If your unit is within one traffic class of its capacity, PacketShaper will not let you create any more classes. This is due to the possibility that two classes will be created in some circumstances. For example, when you create the first child class for a parent, a Default class automatically gets created.

Add a Matching Rule

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To add a matching rule to an existing class:

1. Click the **Manage** tab.

2. In the left window pane, select the class for which you want to add a rule.

3. Click **class** and select **add matching rule**.

4. Complete the entries in the *New Matching Rule* window to define the matching rule for the class. For descriptions of these fields, see Matching Rules.

5. Click **add rule**.

If the info page has flagged one or more of your classes with the configuration error message *attrib iqosMatchingRule = ??? , Failed to add matching rule to traffic class*, you have exceeded the maximum number of matching rules available on your PacketShaper model. (In the CLI, you can display configuration error messages with the **class show *<tclass>*** command.) To free up resources, you need to remove one or more classes or matching rules. Configuration errors will disappear once the total number of matching rules is less than the unit s limit. If you find that you are consistently exceeding your unit s maximum configuration limits, you should consider upgrading your PacketShaper.

Modify a Matching Rule

To modify an existing matching rule:

1. Click the **Manage** tab.
2. In the left window pane, select the class whose matching rule you want to modify.
3. Scroll down and locate the matching rule you want to edit.
4. Click **edit rule**.
5. Modify the matching rule as needed.
6. When finished, click **apply changes**.

Delete a Matching Rule

To delete an existing matching rule:

1. Click the **manage** tab.
2. In the left window pane, select the class whose matching rule you want to delete.
3. Scroll down and locate the matching rule you want to delete.
4. Click **delete rule**.
5. Click **OK** to confirm the deletion.

Classify URL Categories

PacketShaper offers a number of ways you can classify by URL categories:

Create a Category-Based Class

Auto-Discover Category Classes

Discover Categories Under Web-Based Classes

Create a Class for Multiple Categories

Create a Category-Based Class

You can manually create classes for any of the 70+ categories PacketShaper can classify. Alternatively, you can have PacketShaper auto-discover the categories.

PolicyCenter

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a category-based class:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder (such as Inbound or HTTP) that will be the parent of the class you are about to define.

3. Click **class** and select **add**. The *New Traffic Class* screen appears.

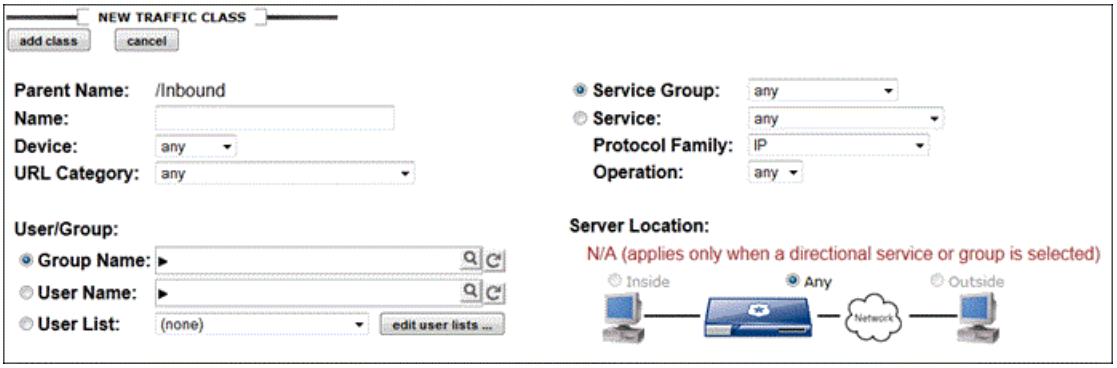

4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).

5. To classify by a particular URL category, choose the **URL Category** from the drop-down list.

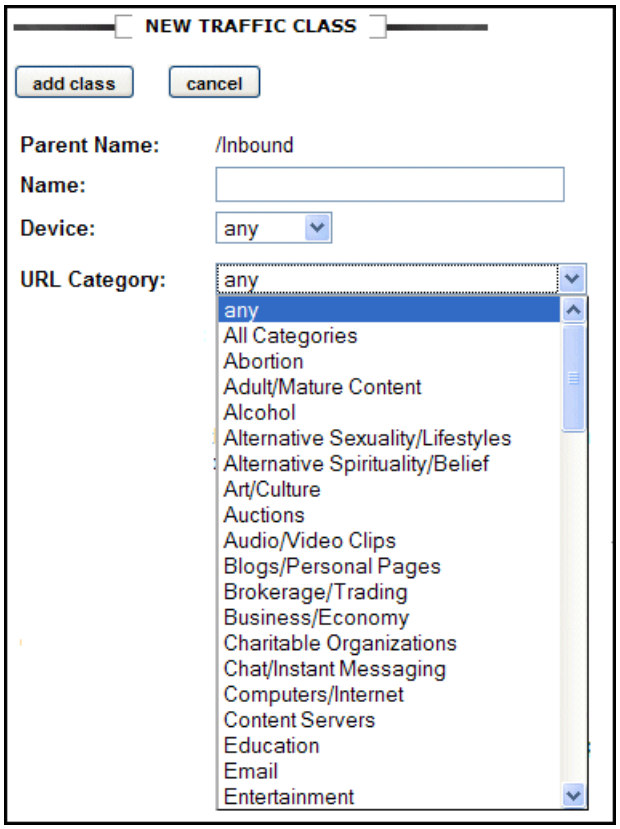

To create a class that discovers child classes for URL categories, choose **All Categories**, and then turn on class discovery.

6. Click **add class**.

7. (Optional) Select **Exception** and click **apply changes**.

> **Note:** When a category class is created in the root, you will usually want to make it an exception class to make sure the web traffic gets classified into the category class instead of in service-based classes (such as HTTP, Facebook, MySpace, or other services that use HTTP as a transport).

8. Make sure WebPulse is enabled.

Auto-Discover Category Classes

To auto-discover category classes, you need to create an "all categories" class, turn on class discovery, and make sure global traffic discovery is enabled. If you also want to auto-discover service-based classes, you should make the "all categories" class an exception so that it appears at the top of the traffic tree and web traffic will get classified into the category class instead of the web-based classes (such as HTTP, SSL, or HTTP-Tunnel).

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To auto-discover category classes:

1. In Inbound, create a class named Categories, and select **All Categories** for **URL Category**.

2. Select **Exception** and enable **Traffic Discovery within Class**. Click **apply changes**.

3. Repeat steps 1 and 2 for Outbound (or copy the class).

4. Make sure WebPulse and global traffic discovery are enabled.

As web traffic appears on your network, PacketShaper will create category classes under the Categories class. For example:

```
Categories
 Business_Economy
 Chat_Instant_Messaging
 Computers_Internet
 Content_Servers
 Email
 Financial_Services
 Shopping
 Social_Networking
 Default
```

Discover Categories under Web-Based Classes

As an alternative to creating an "all categories" class to auto-discover category classes, you can turn on class discovery for web-based classes (such as HTTP, SSL, HTTP-Tunnel). This technique allows you to have different controls for categories depending on the type of web service. For example, you can create different policies for categories that are secure web traffic versus unsecure.

Although you can auto-discover categories for any web-based service that uses HTTP or HTTPS as transport, Symantec recommends that you limit this feature to HTTP, SSL, and HTTP-Tunnel.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To discover categories under web-based classes:

1. If your Inbound traffic tree doesn't already have classes for HTTP, SSL, and HTTP-Tunnel, create classes for these services.

2. Select **Exception** and enable **Traffic Discovery within Class**. Click **apply changes**.

3. Repeat steps 1 and 2 for Outbound (or copy the classes using the **copy class and children** option).

4. Make sure WebPulse and global traffic discovery are enabled.

As web traffic appears on your network, PacketShaper will create category classes under HTTP, SSL, and HTTP-Tunnel. For example:

```
HTTP
 Business_Economy
 Chat_Instant_Messaging
 Computers_Internet
 Content_Servers
 Email
 Financial_Services
 Shopping
 Social_Networking
 Default
SSL
 Brokerage_Trading
 Content_Servers
 Default
```

> **Note:** For category classes to be discovered under an HTTP class, the HTTP class *cannot* have any criteria (such as host, URL, content type, or user agent) defined.

Create a Class for Multiple Categories

If you want to apply the same controls to multiple categories (for instance, apply a never-admit policy to all categories containing adult related content), you can create a single class (Adult_Content) and add matching rules for each category. This technique is also useful to consolidate classes in situations where your PacketShaper is reaching its maximum class limit.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To create a class that matches web traffic from multiple categories:

1. In Inbound, create a class that describes the type of traffic (for example, Adult_Content), and choose the first **URL Category** (such as Pornography) from the drop-down list. Click **add class** .

2. (Optional) Select **Exception** and click **apply changes**.

> **Note:** When a category class is created in the root, you will usually want to make it an exception class to make sure the web traffic gets classified into the category class instead of in service-based classes (such as HTTP, Facebook, MySpace, or other services that use HTTP as a transport).

3. Add a URL category matching rule for each category associated with the class.

4. Repeat steps 1 and 2 for Outbound (or copy the class).

5. Make sure WebPulse is enabled.

Classify Users

PacketShaper has the ability to monitor and control traffic of Active Directory users. You can create classes based on users for the purposes of monitoring and shaping traffic on the network. This functionality allows administrators to monitor users of a service, monitor what services a specific user is using, or manage different policies for a different set of users on the same service.

Create a User-Based Class

Create a User Group Class

Create a Class Based on a User List

Configure PacketShaper as a BCAAA Client

Create a User-Based Class

When PacketShaper is configured as a BCAAA client, you can manually create a class for any of the users in Active Directory.

PolicyCenter

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a user-based class:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder (such as Inbound) that will be the parent of the class you are about to define.

3. Click **class** and select **add**. The *New Traffic Class* screen appears.

4. In the **Name** field, type a class name to identify the user (up to 31 alphanumeric characters; hyphens, underscores, and periods are acceptable; spaces and UTF-8 characters are not allowed).

5. In the **User Name** field, specify the user name using one of the following methods:

- Enter all or any part of the user name (such as the first or the last name) in the **User Name** field. For example, enter *john*, *smith*, or *john.smith*. Then click the Search icon and select the name from the search results list. Note that you may enter UTF-8 characters for the user name (but not for the class name).
- To select the user name from a list all users, leave the **User Name** field empty and click the icon; then select the name from the search results list.

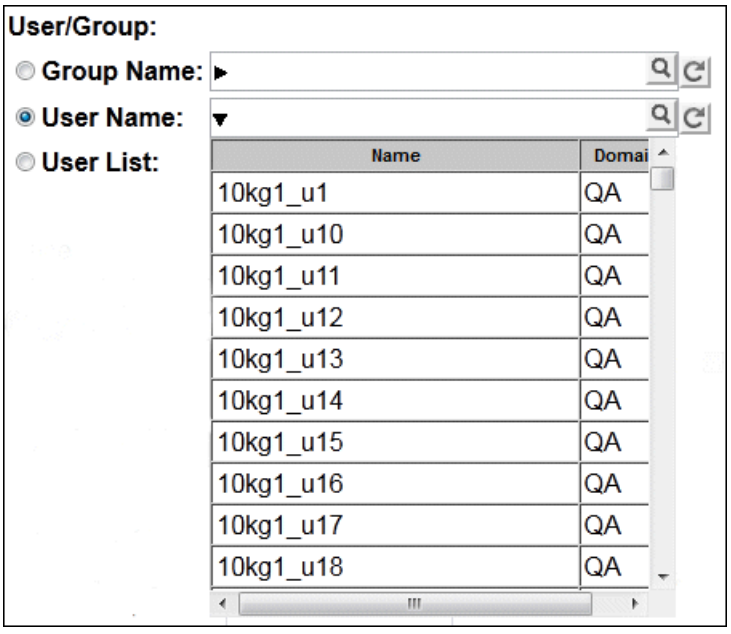

- To list all user names for a specific domain, enter the domain name and click the icon; then select the name from the search results list.

> **Note:** If the search results list is empty, no users were found that match the criteria you entered. Double-check your typing and spelling and try again. Experiment with other types of searches, such as searching by just the first name or just the last name, or display the list of all users and pick the user from the list.

6. Click **add class**.

Additional Information

- When creating classes, PacketShaper does not validate that the specified user or group exists in Active Directory. In addition, PacketShaper allows you to create a child user-based class of a parent group-based class, even if the user isn't in the AD group of the parent class. If you create a class for a user or group that doesn't exist in AD, the class will not have a configuration error, but no traffic will hit the class.
- To clear the **User Name** field of a search string you previously entered, click the button.
- Use the disclosure triangles, and , to expand and collapse the results list.
- When you display the list of all users, these user names are cached locally for the current browser session. To force a refresh of the list (in case users are added, deleted, or modified during your session), click the Refresh button. You will also need to refresh if you change the BCAAA configuration during your current browser session.

Create a User Group Class

When PacketShaper is configured as a BCAAA client, you can manually create a class for any of the user groups in Active Directory.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a class for a user group:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder (such as Inbound) that will be the parent of the class you are about to define.

3. Click **class** and select **add**. The *New Traffic Class* screen appears.

4. In the **Name** field, type a class name to identify the user group (up to 31 alphanumeric characters; hyphens, underscores, and periods are acceptable; spaces and UTF-8 characters are not allowed).

5. In the **Group Name** field, specify the user group name using one of the following methods:

- Enter all or any part of the group name in the **Group Name** field. For example, enter *group-finance or finance*. Then click the Search icon and select the name from the search results list. Note that you may enter UTF-8 characters for the group name (but not for the class name).

- To select the group name from a list all groups, leave the **Group Name** field empty and click the icon; then select the name from the search results list.

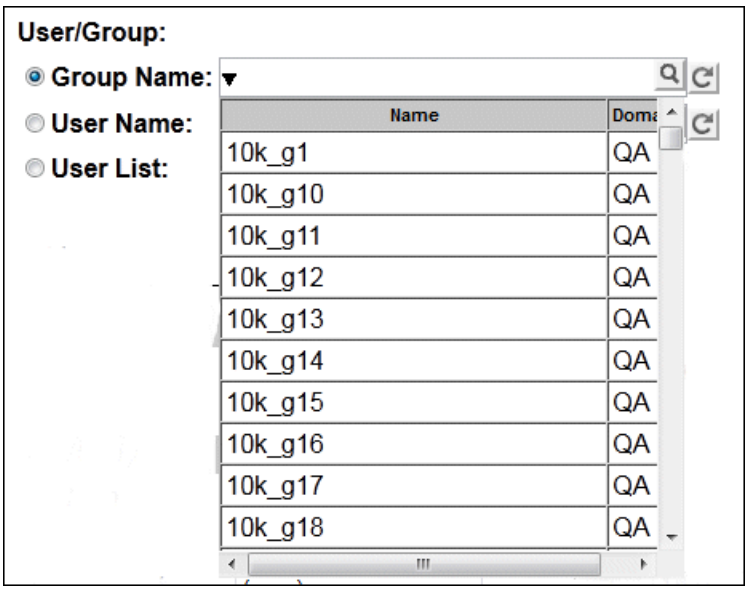

- To list all group names for a specific domain, enter the domain name and click the icon; then select the name from the search results list.

> **Note:** If the search results list is empty, no groups were found that match the criteria you entered. Double-check your typing and spelling and try again. Experiment with other types of searches, or display the list of all groups and pick the group from the list.

6. Click **add class**.

Additional Information

- When creating classes, PacketShaper does not validate that the specified user or group exists in Active Directory. In addition, PacketShaper allows you to create a child user-based class of a parent group-based class, even if the user isn't in the AD group of the parent class. If you create a class for a user or group that doesn't exist in AD, the class will not have a configuration error, but no traffic will hit the class.
- Each PacketShaper model has a limit to the number of group-based classes that can be created. (See Configuration Limits in PacketGuide.) Once you have exceeded the maximum number of user groups that can be used in matching rules, you must delete an existing group-based class before you can create a new one.
- To clear the **Group Name** field of a search string you previously entered, click the button.
- Use the disclosure triangles, and , to expand and collapse the results list.
- When you display the list of all groups, these group names are cached locally for the current browser session. To force a refresh of the list (in case groups have been added, deleted, or modified during your session), click the Refresh button. You will also need to refresh if you change the BCAAA configuration during your current browser session.

Create a Class Based on a User List

When PacketShaper is configured as a BCAAA client, you can create a named list of users and groups, and then create a traffic class based on this user list. This allows you to monitor and control traffic for a set of users and groups.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a class for a user list:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder (such as Inbound) that will be the parent of the class you are about to define.

3. Click **class** and select **add**. The *New Traffic Class* screen appears.

4. In the **Name** field, enter a descriptive name for the list (up to 31 alphanumeric characters; hyphens, underscores, and periods are acceptable; spaces and UTF-8 characters are not allowed).

5. Select **User List**.

6. Select an existing list from the **User List** drop-down.

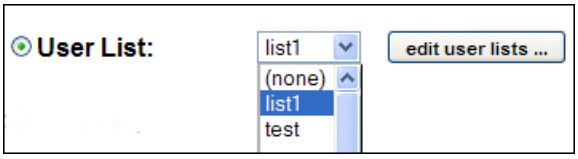

or

Click **edit user lists** to create or modify a list. The User List Editor opens.

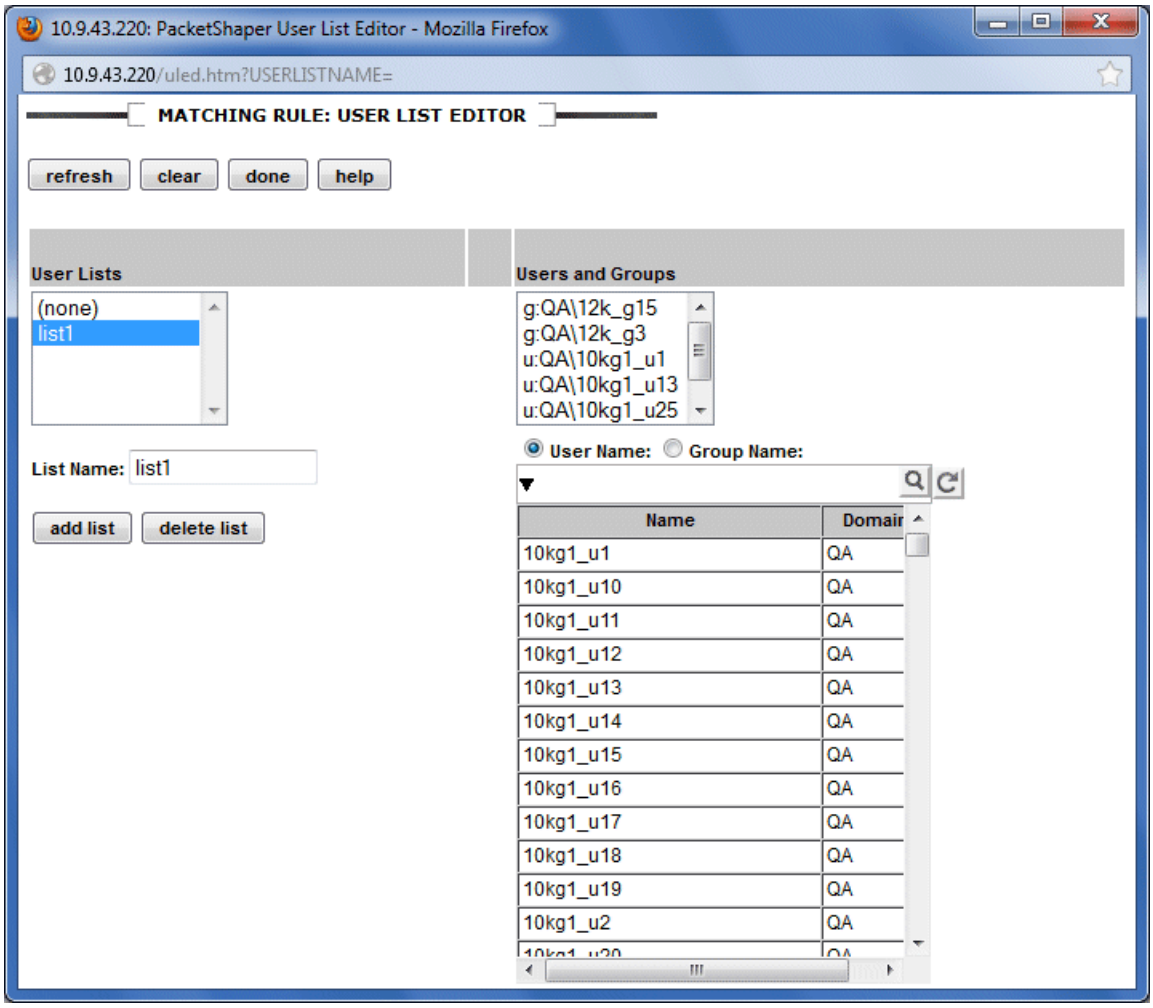

7. After you have selected the user list, click **add class**.

Open the User List Editor

The Advanced UI offers a User List Editor for creating and managing user lists. You can access the User List Editor when creating or modifying classes or matching rules.

PolicyCenter To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To open the User List Editor:

1. While creating/modifying a class or a matching rule, select **User List** and then click **edit user list**. The User List Editor appears.

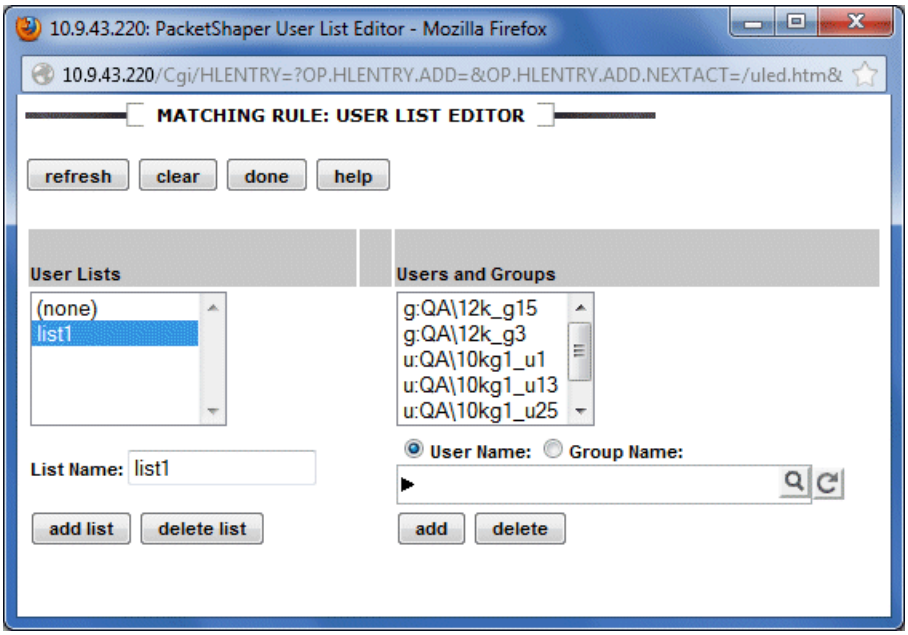

2. You can now create, modify, or delete user lists.

Create User Lists

To create a class for a custom set of users or groups, you can create a *user list*. A user list can contain a combination of user names and group names. The User List Editor is only accessible when creating or modifying classes or matching rules. Thus, to create a user list, you must first create a class or matching rule.

PolicyCenter To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To create a user list:

1. Open the User List Editor.

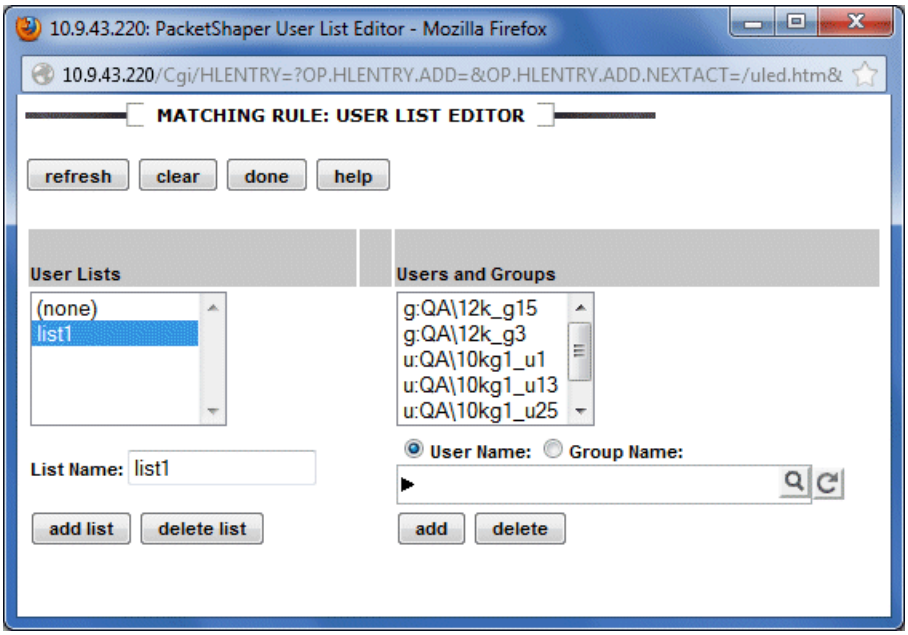

2. In the *User List Editor* window, enter a descriptive name for the list in the **List Name** field.

3. Click **add list** to generate the new user list, and click OK.

4. Select **User Name** or **Group Name**, the use any of the following methods to locate the name:

   - Click the Search icon to display the complete list of user or group names. Use the scroll bar to locate the desired name on the search results list. You may enter UTF-8 characters for the user or group name.

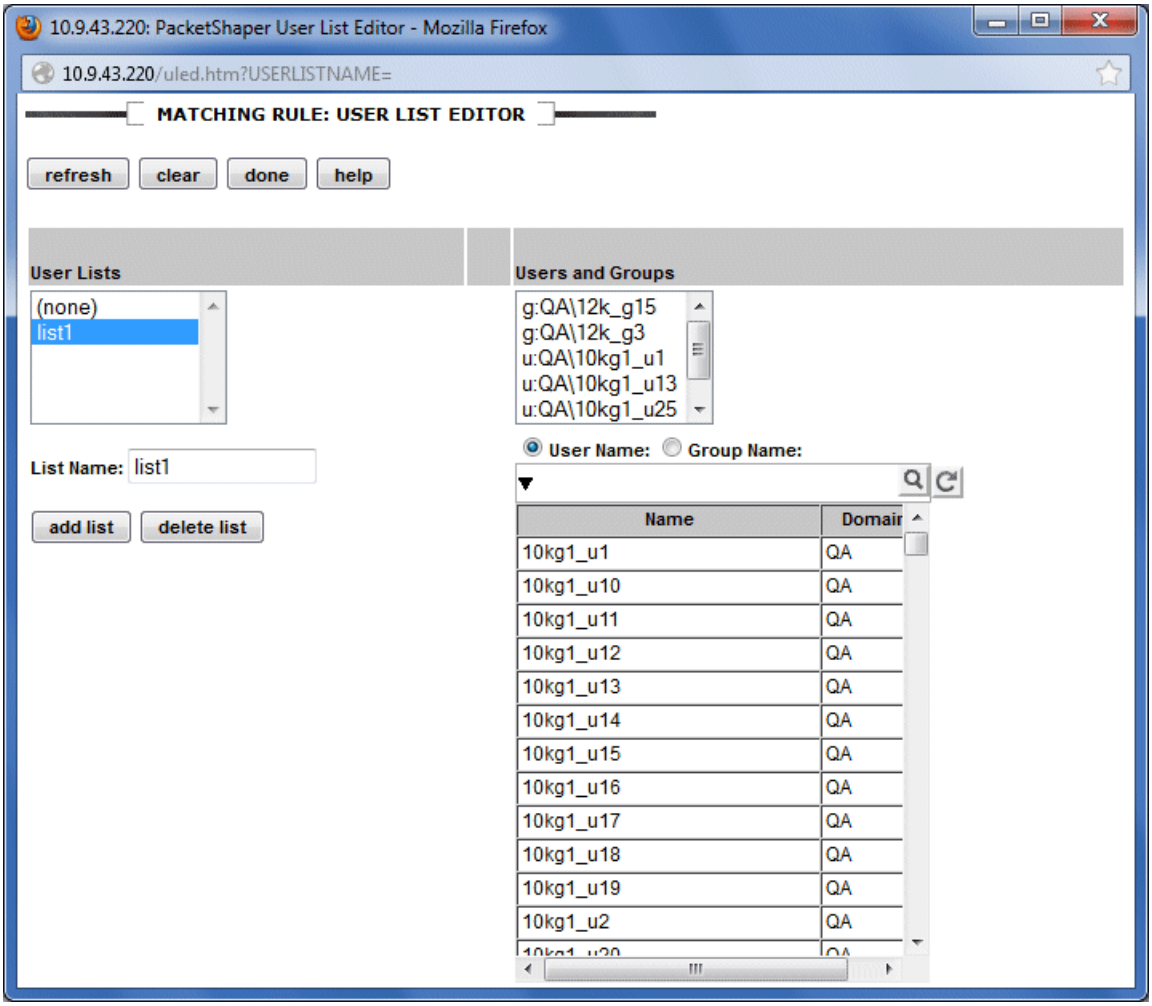

- Enter all or any part of the name. For example, for a user name enter *john*, *smith*, or *john.smith*. Or for a group name, enter *group-engineering* or *engineering*. Then click 🔍 and select the name from the search results list.
- To list all names for a specific domain, enter the domain name and click 🔍 ; then select the name from the search results list.

> **Note:** If the search results list is empty, no users or groups were found that match the criteria you entered. Double-check your typing and spelling and try again. Experiment with other types of searches, such as searching by just the first name or just the last name, or display the list of all users/groups and pick the name from the list.

5. Click **add**. The name is added to the **Users and Groups** box. Note that the system automatically adds **u:** before a user name and a **g:** before each group name.

6. Repeat steps 4 and 5 to continue adding users and/or groups and then click **done** to return to the *Matching Rule* or *Traffic Class* window.

7. To include the selected user list in the matching rule definition, click **add class** (when creating a new class) or **apply changes** (when modifying an existing class).

Edit a User List

You can use the User List Editor to add user and group names to an existing user list, or delete names from the list.

PolicyCenter To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To modify a user list:

1. Open the User List Editor.

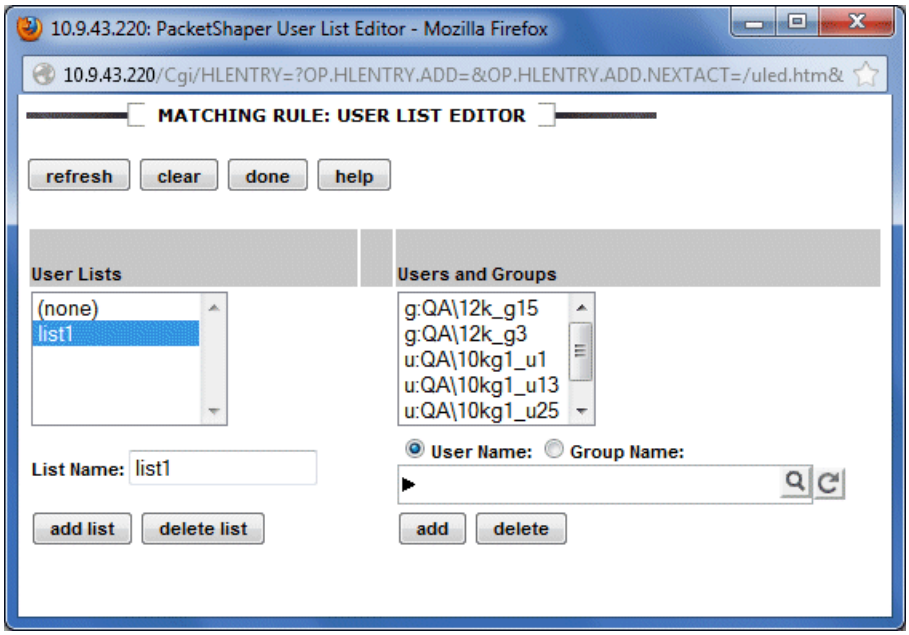

2. In the *User Lists* box on the left, select the name of the list you want to edit. The user and group names associated with this list are displayed in the *Users and Groups* box on the right.

3. To add another user or group to this list, select **User Name** or **Group Name**, then use any of the following methods to locate a name:

- Click the Search icon to display the complete list of user or group names. Use the scroll bar to locate the desired name on the search results list.
- Enter all or any part of the name (such as the first or the last name) in the **User Name** field. For example, for a user name, enter *john*, *smith*, or *john.smith*. Or for a group name, enter *group-support* or *support*. Then click and select the name from the search results list. You may enter UTF-8 characters for the user or group name.
- To list all user or group names for a specific domain, enter the domain name and click ; then select the name from the search results list.

> **Note:** If the search results list is empty, no users or groups were found that match the

> criteria you entered. Double-check your typing and spelling and try again. Experiment with other types of searches, such as searching by just the first name or just the last name, or display the list of all users/groups and pick the name from the list.

4. Click **add**. The name is added to the **Users and Groups** box. Note that the system automatically adds **u:** before a user name and a **g:** before each group name.

5. To remove a user or group from this list, select the name in the *Users and Groups* box and click **delete**, then click OK to confirm. To select multiple users for deletion, hold down Ctrl and click each name or use the Shift-click technique to select consecutive user names.

6. Click **done** to return to the *Matching Rule* or *Traffic Class* window.

7. Click **apply changes**.

Delete a User List

You can use the User List Editor to delete a user list that is no longer needed. Note that the Editor prevents you from deleting a user list that is currently specified in a matching rule.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To delete a user list:

1. Open the User List Editor.

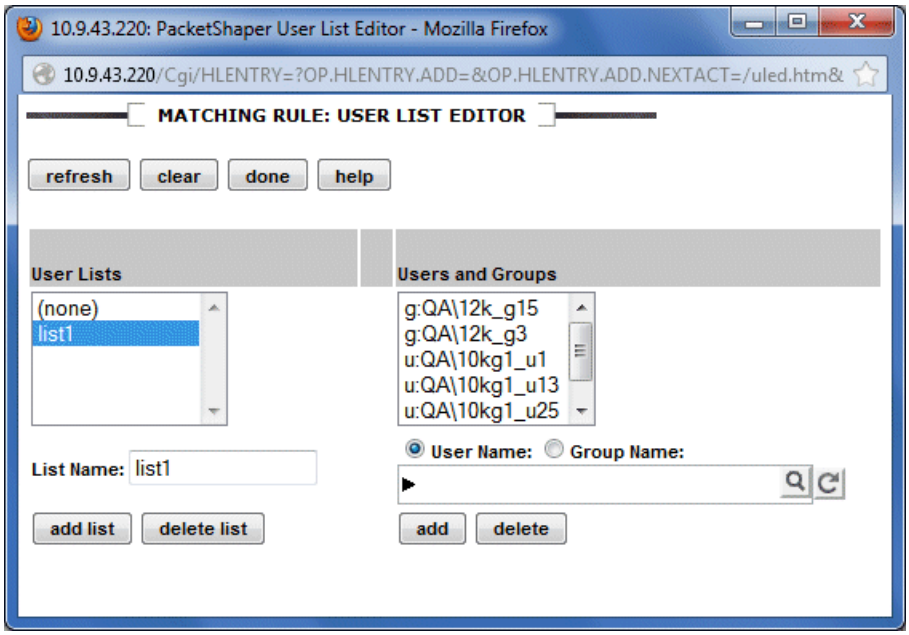

2. In the *User Lists* box on the left, select the name of the list you want to delete. The users and groups associated with this list are displayed in the *Users and Groups* box on the right.

3. Click **delete list** and click OK to confirm.

4. Click **done** to return to the *Matching Rule* or *Traffic Class* window.

5. Click **apply changes**.

> **Note:** If you try to delete a user list that is currently specified in a matching rule, you will see the following message: *User list xxxx is in use*. Before you can delete this user list, you need to determine which classes specify this user list and remove this matching rule.

Manually Create an IPv6 Class in the Advanced UI

The PacketShaper can auto-discover a number of classes based on IPv6 services–such as HTTP-IPv6, FTP-IPv6, and CIFS-IPv6–but you can manually create these classes as well.

When defining an IPv6 class in the Advanced UI, you must specify the IPv6 protocol family in order to see the list of IPv6-based services.

PolicyCenter

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

1. Click the **Manage** tab.
2. In the left window pane, select the class or folder that will be the parent of the class you are about to define.
3. Click **class** and select **add**. The *New Traffic Class* screen appears.
4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).
5. For the **Protocol Family**, select **IPv6**; once you do this, the **Service** list (step 6) then displays the IPv6-based services.

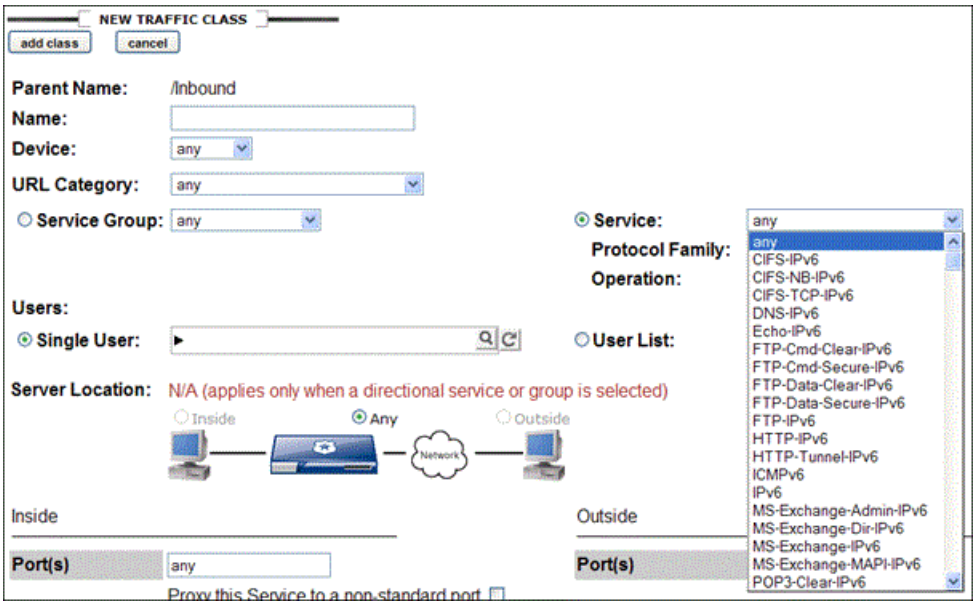

6. Select the desired IPv6 service from the **Service** list.
7. Click **add class**.

Create a Class for an IPv6 Server

The PacketShaper can auto-discover a number of classes based IPv6 services—such as HTTP-IPv6, FTP-IPv6, and CIFS-IPv6—but you can manually create these classes as well.

If you have a custom IPv6 application, you can create a class for the server's IPv6 address. Alternatively, you can create a class for an IPv6 subnet.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

1. Create a class, choosing **IPv6** for the **Protocol Family**.

2. Enter the IPv6 address in the **IP Address** field (for Inside or Outside, depending on where the server is located).

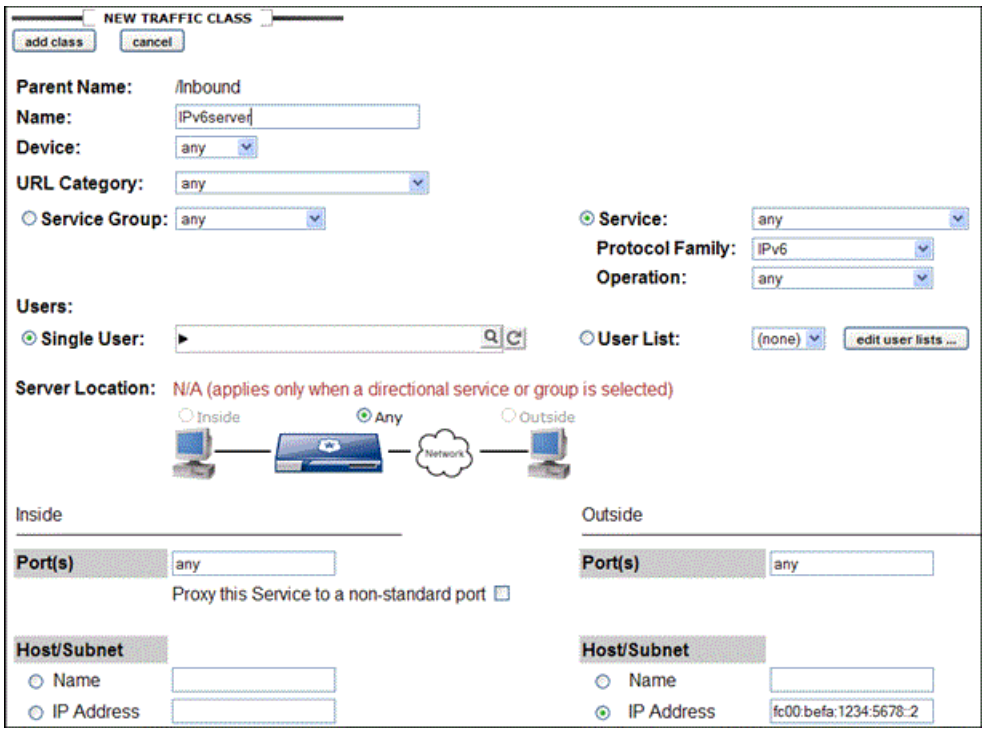

3. *(Optional)* Enable **Traffic Discovery within Class** if you want IPv6 services to be auto-discovered as children of the IPv6 server/subnet class.

Create a Service-Centric Traffic Tree for IPv6

In the following traffic tree example, there is a parent class that classifies all FTP traffic (IPv4 and IPv6), with child classes that break out the IPv4 and IPv6 flows. It allows you to track metrics for all FTP as well as for IPv4 and IPv6 FTP traffic.

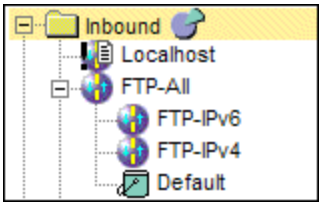

The FTP-All class has a set of matching rules for the FTP service (which classifies IPv4 FTP traffic) and another set of matching rules for the FTP-IPv6 service.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To create a combined IPv4+IPv6 class with children:

1. In the Inbound root, create an "All" class for the IPv4-based service.

2. Add a matching rule to this class for the IPv6 version of this same service.

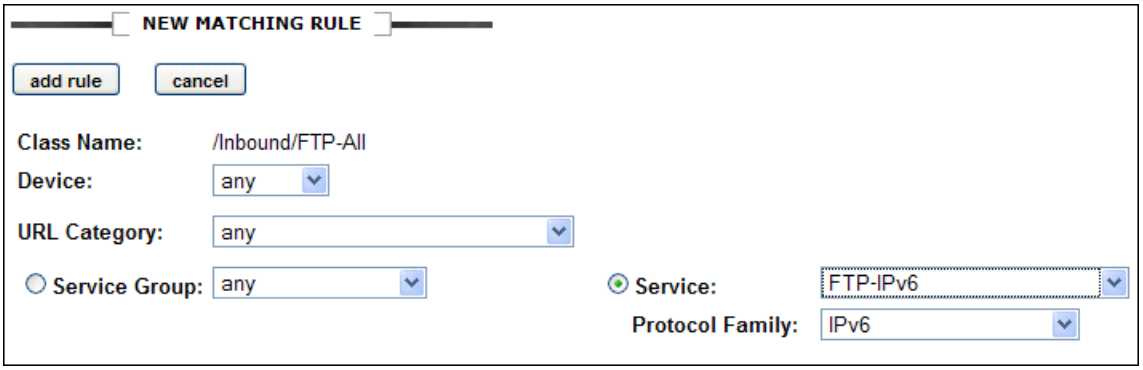

3. Create a child class of the "All" parent that classifies the IPv4-based service.

4. Create a child class of the "All" parent that classifies the IPv6 version of this same service.

5. Copy the parent and child classes to the Outbound root.

Create an IPv6-Based Branch in the Traffic Tree

In this traffic tree example, there is an IPv6 parent class and child classes for each IPv6 service.

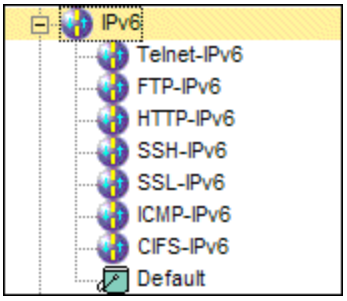

You can either auto-discover the child IPv6 service classes or you can manually create them.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

1. In the Inbound root, create an IPv6 class (service=IPv6) in Advanced UI or Sky UI.

2. Turn on **Traffic Discovery within Class** for the IPv6 class.

or

For each IPv6 service on your network, manually create a child class of the IPv6 class.

3. Copy the parent and child classes to the Outbound root. (Advanced UI)

> **Note:** This step is not necessary for Sky UI because the class is automatically created in Inbound and Outbound.

Create a Site-Based Traffic Tree for IPv6 Subnets

When building a site-based class tree in an IPv6 environment, you can create a class for each site by defining the applicable IPv6 subnets as class matching rules.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

1. Create a class, choosing IPv6 for the **Protocol Family**.

2. To define an IPv6 subnet, enter the prefix (for example, *2001:db8:420::*) in the **Subnet** field.

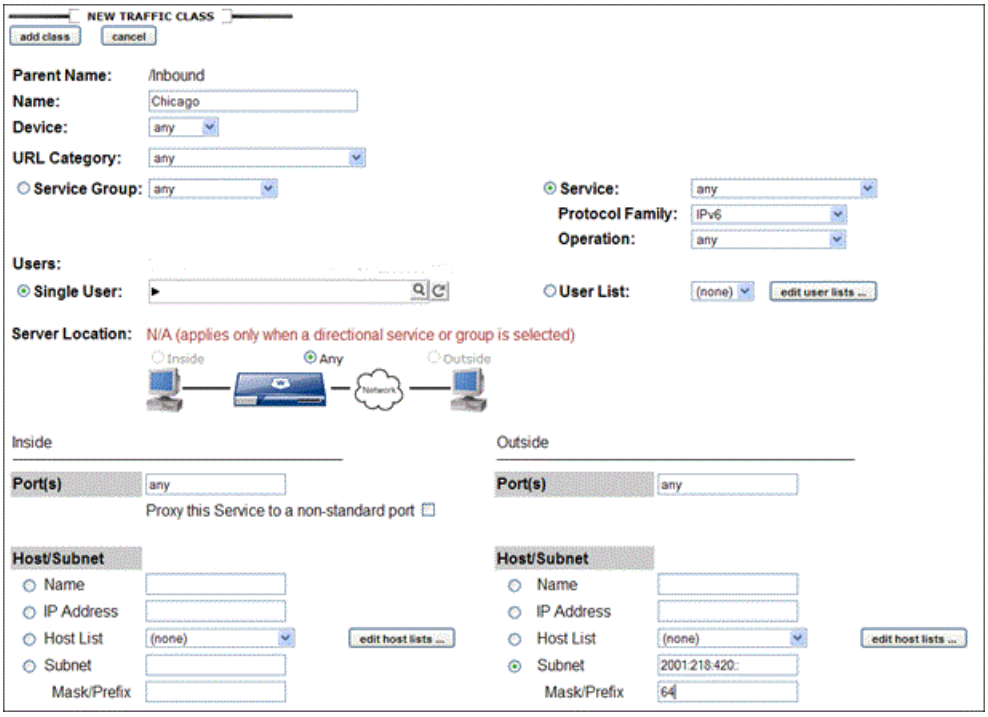

3. Specify the prefix length (for example, *48*) in the **Mask/Prefix** field.

4. (Optional) **Enable Traffic Discovery within Class** if you want IPv6 services to be auto-discovered as children of the IPv6 subnet class.

Create a Citrix-ICA Class

If traffic discovery is enabled, PacketShaper will automatically discover Citrix traffic and create classes for published applications. For instance, when PacketShaper detects Citrix traffic, a Citrix class will be created, and class discovery will be enabled on that class. Because class discovery is on, PacketShaper will also create a Citrix/Default class and Citrix child classes for any published applications.

If traffic discovery is disabled, you can create Citrix classes manually.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To manually create a Citrix-ICA class:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder that will be the parent of the class you are about to define.

3. Click **class** and select **add**.

4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).

5. For the **Service**, select **Citrix-ICA**.

6. *Optional:* To specify a Citrix-ICA criterion, scroll down to the **Criterion** field, and select **Published Application**, **Client Name** or **Priority** from the drop-down list. **show screen...**

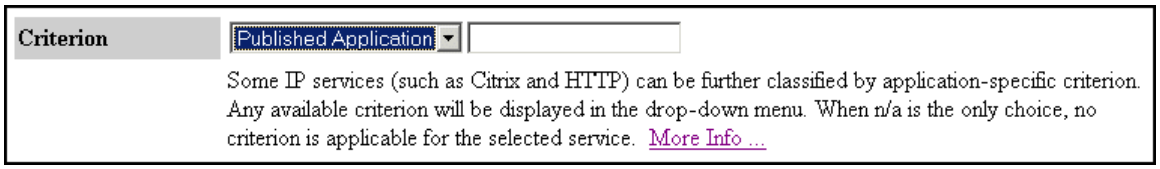

See Citrix Criteria for details.

7. Type the criterion (name or priority) in the empty text field.

6. Click **add class**.

For more information about the other entries in the *New Traffic Class* window, see Matching Rules. A matching rule can contain only one application-specific criterion. To combine multiple criteria, use two traffic classes, a parent and a child.

Additional Information

- In order to classify Citrix published applications or priority bit tags when using MetaFrame Presentation Server 3.0+, the Citrix Session Reliability feature must be disabled.
- To control Citrix traffic after you have created Citrix-based classes, see Manage Citrix Performance.

Citrix Criteria

PacketShaper offers three ways to classify Citrix-ICA traffic:

- by published application
- by client name
- by priority

Classifying by Published Application

> **Note:** PacketShaper can automatically discover published application names when traffic discovery is enabled or when Citrix has traffic discovery enabled within the class. For instance, when PacketShaper detects Citrix-ICA traffic from the application published as "Microsoft Word," it automatically creates a Citrix-Microsoft_Word class.

In order to classify Citrix published applications when using MetaFrame Presentation Server 3.0+, the Citrix Session Reliability feature must be disabled.

To manually create a class for a Citrix-ICA published application, specify Citrix-ICA as the service, and select **published application** from the **Criterion** drop-down list. **show screen...**

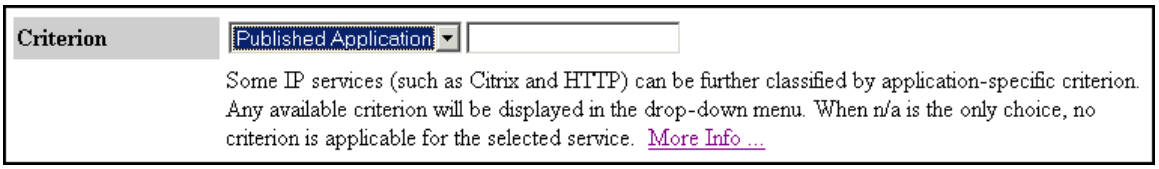

You then must supply a specific name (up to 99 characters). The published application name you enter must exactly match the name configured on the Citrix server and client.

For example, you could create a class for Citrix-ICA traffic that carries the published application PeopleSoft. To create a PeopleSoft class as a child class of the Citrix class, you would specify the following attributes:

- Name: PeopleSoft
- Protocol: IP
- Service: Citrix-ICA
- Location: any
- Criterion: Published Application
- Criterion (second field): PeopleSoft

> **Note:** The entry in this second Criterion field must match the name for the published application defined in the Citrix Program Neighborhood.

Classifying by Client Name

To classify by a Citrix-ICA client name, specify Citrix-ICA as the service, and select **client name** from the **Criterion** drop-down list. **show screen...**

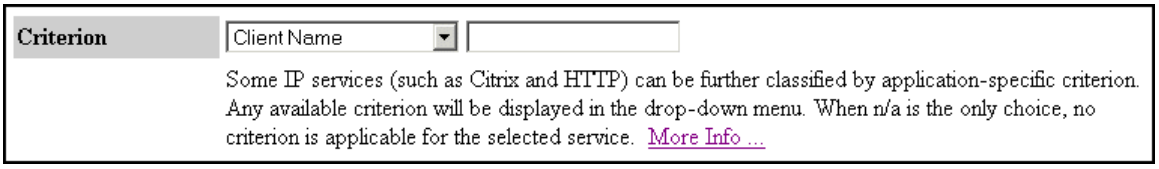

You then must supply a specific name (up to 99 characters). The client name you enter must exactly match the name configured on the Windows client.

Classifying by Priority

Citrix-ICA has the ability to assign different priority levels to virtual channels within a single ICA flow, and PacketShaper can classify the traffic according to these priority levels. With Citrix-ICA specified as the service, select **priority** from the **Criterion** drop-down list. **show screen...**

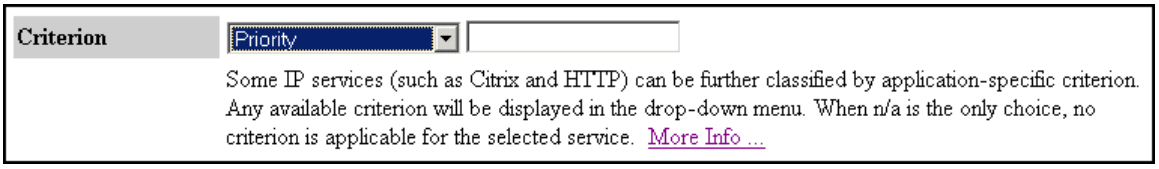

Supply a value from 0 to 3 (with 0 being the highest priority). You may want to create Citrix child classes for each priority level. Note that the Citrix priority tag of 3 is always print traffic.

In order to classify priority bit tags when using MetaFrame Presentation Server 3.0+, the Citrix Session Reliability feature must be disabled.

> **Note:** The Citrix priority numbers are used for classification purposes only. Once the traffic is categorized by Citrix priority, you can use PacketShaper policies to manage bandwidth on the class.

Classify Web Traffic

PacketShaper can do much more than identify port 80 traffic. You can create web classes based on a specific IP address or host name as well as a URL to a specific web page. By integrating with the WebPulse cloud service, you can create classes based on URL categories or web applications and operations.

Create Web-Based Classes

Create a Category-Based Class

Classify Web Applications

List Web Applications

Configure WebPulse

Create Web Classes

Because PacketShaper can identify specific attributes of web traffic, you can create classes based on these attributes. You can create web classes based on a specific IP address or host name as well as a URL to a specific web page.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a web class:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder that will be the parent of the class you are about to define. Usually, you will want web classes to be children of the HTTP class.

3. Click **class** and select **add**.

4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).

5. For the **Service**, select **HTTP (Web)** or **HTTP-IPv6**.

> **Note:** You will need to choose **IPv6** for the **Protocol Family** if you want to select the HTTP-IPv6 service.

6. Scroll down to the **Criterion** field, and select **Host DNS Name or IP Address** or **URL** from the drop-down list. **show screen...**

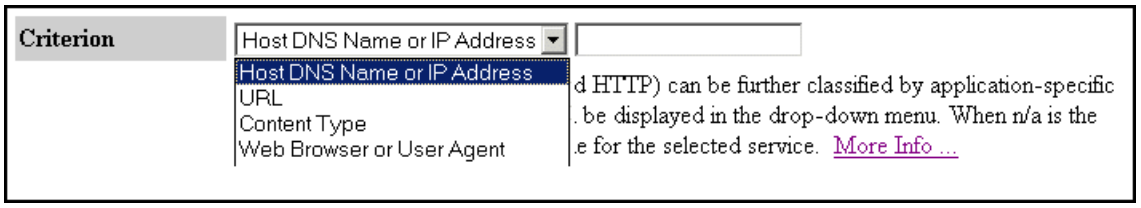

> **Note:** You cannot enter an IPv6 address as a web criterion.

7. If you selected **Host DNS Name** for the criterion, enter the DNS name (for example, www.test.com) or IP address (such as 10.10.10.10) of the website in the empty text field.

or

If you selected **URL**:

- In the empty text field next to URL, type only the portion of the URL that follows the server hostname (http://hostname). For example, to specify a remote website *http://www.mycompany.com/FavoriteStuff*, enter **/FavoriteStuff** in the URL **Criterion** field. See URL Syntax for details.

- Enter the hostname in the **Name** field (in the Inside or Outside column in the window, depending on where the web server resides). For example, enter **www.mycompany.com** into the **Name** field of the Outside column.

8. Click **add class**.

For more information about the other entries in the *New Traffic Class* window, see Matching Rules.

Additional Information

- To classify by DNS name, a DNS server must be configured and that server must be able to resolve the name of the site.
- A matching rule can contain only one application-specific criterion. To combine multiple criteria, use two traffic classes, a parent and a child.

Classify Web Applications

By using the WebPulse cloud service, PacketShaper can expand its classification capabilities to include over 100 web applications, such as Facebook, Gmail, LinkedIn, Twitter, and Flickr. In addition to classifying by web application, you can classify by specific web operations (such as Download Attachment, Play Video, and Send Email). Thus, you can create the following types of classes:

- A specific web application for any operation (Service=Facebook, Operation=any)
- A specific operation for any web application (Operation=Upload Pictures, Service=any)
- A specific web application for a specific operation (Service=Facebook, Operation=Upload Pictures)

You can manually create classes based on web applications and operations, or you can have PacketShaper autodiscover the web applications.

Manually Creating Web Application Classes

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a web application class:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder (such as Inbound) that will be the parent of the class you are about to define.

3. Click **class** and select **add**.

4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).

5. To classify by a particular web application, choose the **Service** from the drop-down list.

6. *(Optional)* To classify by a specific web operation (such as Download Attachment), choose the **Operation** from the drop-down list. If you selected a web application for the Service in the previous step, the **Operation** drop-down will list only the operations that are applicable to that application.

7. Click **add class**.

8. Make sure WebPulse is enabled.

Autodiscovering Web Applications

Assuming traffic discovery and WebPulse are both enabled, PacketShaper automatically discovers web applications and creates classes for them in the traffic tree.

To autodiscover web applications:

1. Make sure global traffic discovery is enabled.

2. Make sure WebPulse is enabled.

Additional Information

- Although you can create web application classes when WebPulse is disabled, there won't be any classification into these classes until you enable WebPulse.
- Web operations cannot be autodiscovered.

View Lists of Web Applications and Operations

To determine which web applications and operations WebPulse can identify, go to the WebPulse setup screen. You can create traffic classes for any of these applications and operations, or if traffic discovery and WebPulse are both enabled, PacketShaper will auto-discover these applications.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

1. Click the **Setup** tab.

2. From the list of setup pages, select **WebPulse**. The WebPulse setup page appears.

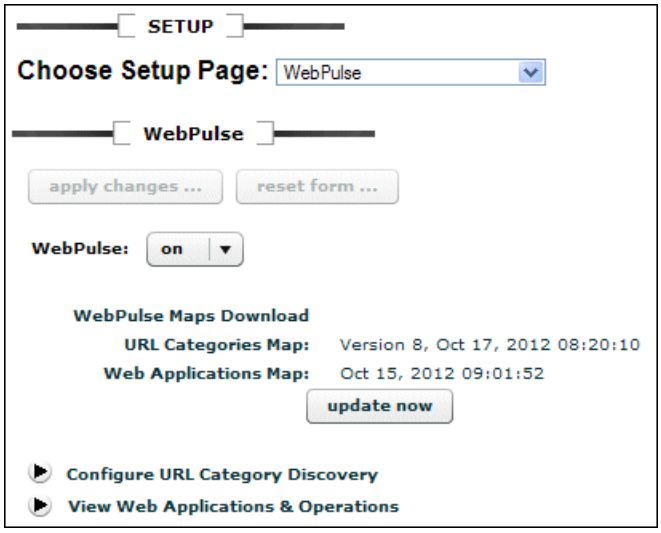

3. Click the next to **View Web Applications & Operations**. The Web Applications are listed.

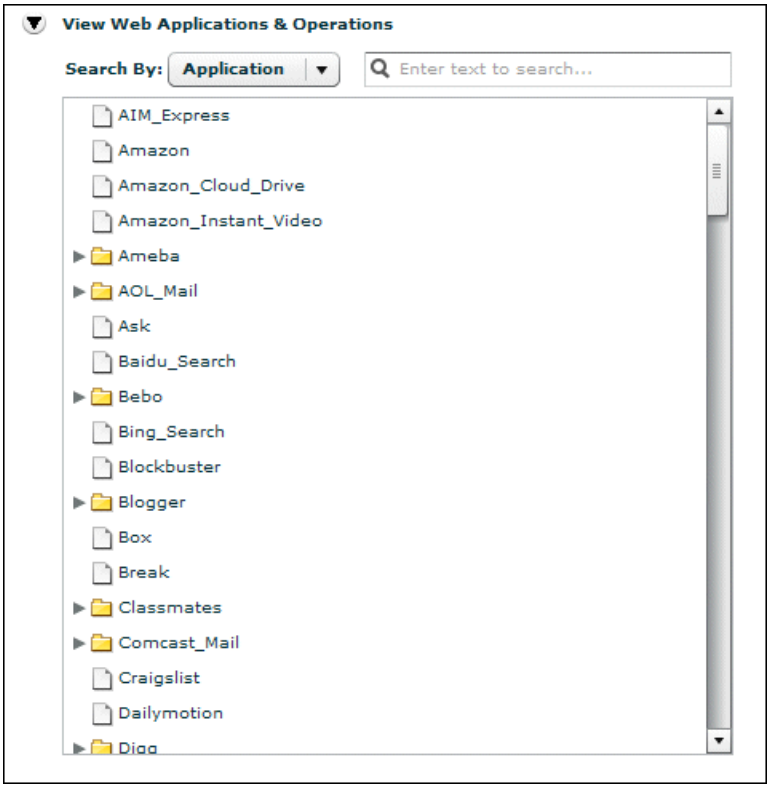

4. To see which operations are applicable to an application, click the next to the application name. The applicable operations are listed below the application.

5. To list the applications that are assoicated with a particular operation:

   a. For **Search by**, select **Operation**.

   b. In the text box, type the operation (such as *upload*). The list is then filtered by the specified operation, displaying only the applications that can be classified by that operation.

   c. To remove the search filter, click .

6. To search for a specific application:

   a. For **Search by**, select **Application**.

   b. In the text box, type the application (such as *LinkedIn*). Only the specified application is displayed in the list.

   c. To remove the search filter, click .

Additional Information

- Web applications that WebPulse can identify are also listed on the **Services** drop-down on the New Traffic Class page; they are mixed in with other PacketShaper classification services. See "Classify Web Applications" (Advanced UI, Sky).
- Operations are also listed on the **Operations** drop-down when creating a class.

Create FTP Classes

In order to classify and put controls on specific types of FTP downloads, you can create an FTP class that matches FTP traffic that has a specific file extension or filename. The class must be based on the FTP-Data-Clear or FTP-Data-Clear-IPv6 service.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create an FTP class:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder that will be the parent of the class you are about to define. Usually, you will want it to be a child of the FTP class.

3. Click **class** and select **add**.

4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).

5. For the **Service**, select **FTP-Data-Clear** or **FTP-Data-Clear-IPv6**.

> **Note:** You will need to choose **IPv6** for the **Protocol Family** if you want to select the FTP-Data-Clear-IPv6 service.

6. Scroll down to the **Criterion** field, and make sure **File Name** is selected. **show screen...**

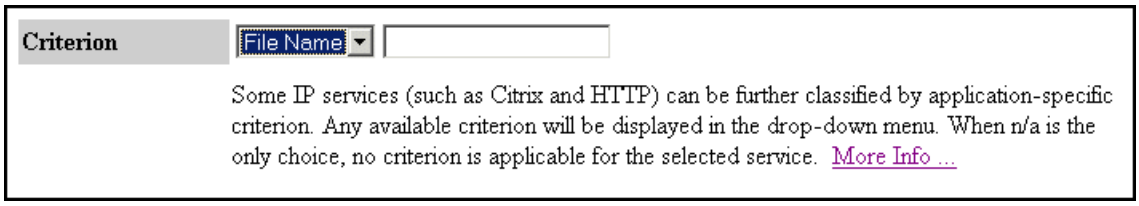

7. Enter the filename or file extension in the empty text field for example, **.mp3*.

8. Click **add class**.

For more information about the other entries in the *New Traffic Class* window, see Matching Rules.

> **Note:** A matching rule can contain only one application-specific criterion. To combine multiple criteria, use two traffic classes, a parent and a child.

Create a Database Class

PacketShaper can classify traffic by a specific Oracle or PostgreSQL database name. The database classes can be created either automatically or manually.

Automatic Creation of Database Classes

When traffic discovery is enabled, PacketShaper will automatically discover Oracle or PostgreSQL database names and create classes for them. For instance, when PacketShaper detects Oracle traffic, an Oracle class will be created, and class discovery will be enabled on that class. Because class discovery is on, PacketShaper will also create an Oracle/Default class and Oracle child classes for any Oracle databases.

If your traffic tree already has an Oracle or PostgreSQL class, you can manually enable class discovery, whereupon the database-named classes will be automatically created as child classes of Oracle or PostgreSQL.

Manual Creation of Database Classes

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a class for an Oracle or PostgreSQL database:

1. Click the **Manage** tab.
2. In the left window pane, select the class or folder that will be the parent of the class you are about to define.
3. Click **class** and select **add**.
4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).
5. For the **Service**, select **Oracle-netv2** or **PostgreSQL**.
6. Scroll down to the bottom of the page, and make sure **Database Name** appears in the **Criterion** field. **show screen...**

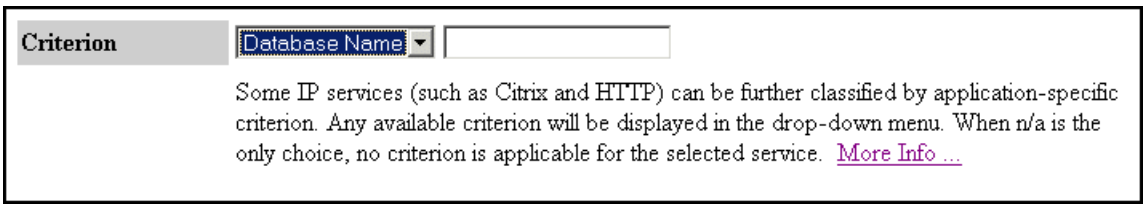

7. Type the database name in the empty text field.
8. Click **add class**.

For more information about the other entries in the *New Traffic Class* window, see Matching Rules.

Note: A matching rule can contain only one application-specific criterion. To combine multiple criteria, use two traffic classes, a parent and a child.

Classify Traffic Based on VLAN ID or Priority

PacketShaper has the ability to classify application traffic within multiple virtual LANs on a single Ethernet segment (an 802.1Q trunk) and ISL (Inter-Switch Link, a Cisco protocol), as well as classify by VLAN priority. You can create a VLAN identification matching rule that is based on a single VLAN ID or priority, a range of IDs or priorities, or any ID or priority.

For a multiple VLAN environment (where packets passing through the unit will have more than one VLAN header), only the outermost VLAN header is used for classification. The other headers are ignored. If packets will be tagged with multiple VLAN headers, be sure to verify that the additional header information does not create a final packet size greater than 1548 bytes before the packet reaches the PacketShaper. (PacketShaper can accept packets up to 1548 bytes in size and will drop any packets that exceed this size.)

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

VLAN Identification (802.1Q or ISL)

To set up a VLAN identification matching rule:

1. When creating a class or adding a matching rule, click **VLAN**. The *Matching Rule: VLAN* window opens. **show screen...**

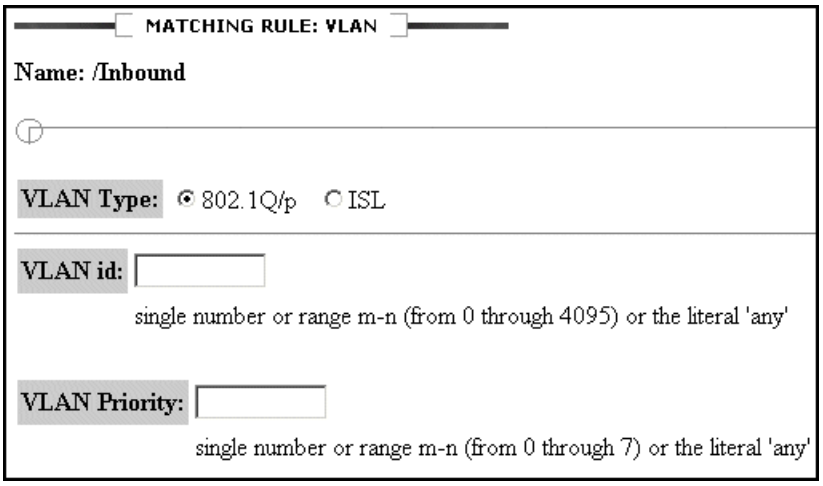

2. Select the VLAN type: **802.1Q/p** or **ISL**.

3. In the **VLAN id** field, enter one of the following:

   - To specify a single VLAN ID, enter a number between 0 4095 for 802.1Q or between 0 32767 for ISL.
   - To specify a range of VLAN IDs, enter the first and last numbers, separated by a dash (for example, 100-200)

- To match any traffic flow that has a VLAN ID, type **any**.

4. Click **OK**.

5. Click **add class** or **add rule**.

> **Note:** When creating a child class under a VLAN ID parent, if you specify **any** in the VLAN ID field, you will get an error message. You must either leave the VLAN ID field blank or specify the same VLAN ID as the parent class has. This is not a restriction in the command-line interface.

VLAN Priority (802.1p)

To set up a VLAN priority matching rule:

1. When creating a class or adding a matching rule, click **VLAN**. The *Matching Rule: VLAN* window opens. **show screen...**

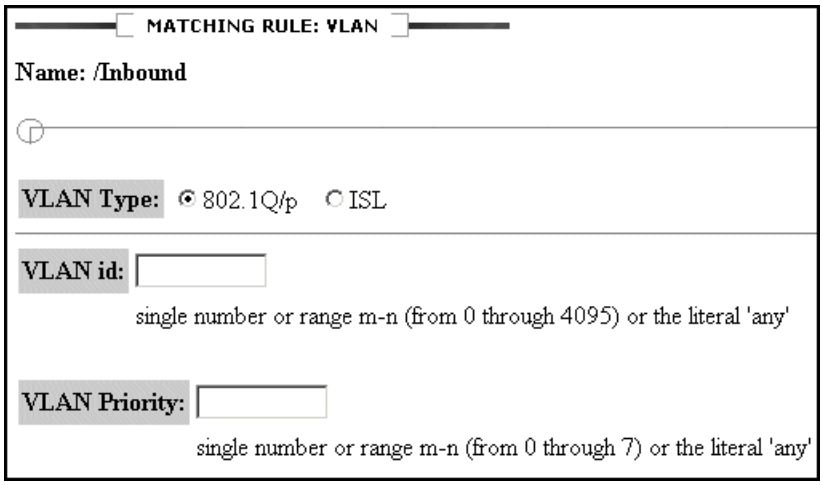

2. Select **802.1Q/p** for the VLAN type.

3. In the **VLAN Priority** field, enter one of the following:

- To specify a single VLAN priority, enter a number between 0 7.
- To specify a range of VLAN priorities, enter the first and last numbers, separated by a dash (for example, 1-3).
- To match any traffic flow that has a VLAN priority tag, type **any**.

4. Click **OK**.

5. Click **add class** or **add rule**.

Once the VLAN class is created, you can enable discovery within the class to discover all applications running over this VLAN that have the ID or priority you specified in the matching rule.

Classify Diffserv Code Point Traffic

Network routers can apply differentiated grades of service to packet streams based on a Diffserv Code Point (DSCP) marker in an IP packet header. PacketShaper is able to create classes based on the DSCP marker.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To set up a Diffserv Code Point matching rule:

1. When creating a class or adding a matching rule, click **diffserv**. The *Matching Rule: Diffserv* window opens. **show screen...**

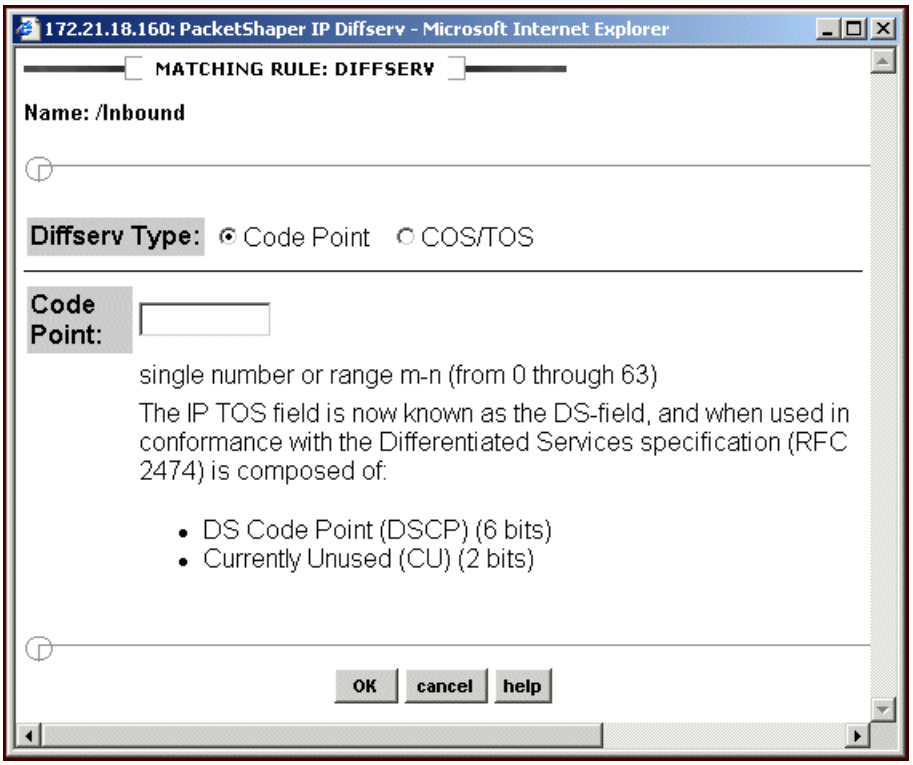

2. Select **Code Point** for the Diffserv Type.

3. Enter the **Code Point** value (0 63) or a range of values.

4. Click **OK**.

5. Click **add class** or **add rule**.

Classify Diffserv COS/TOS Traffic

Many routing protocols make routing decisions based on the Type of Service (TOS) field in the IP header. This field is an indicator of the quality of service that is expected. Applications can set this field to tell routers how to prioritize packets. For example, weighted fair queuing (WFQ) algorithms in routers use this information.

The TOS field contains two subfields: IP Precedence (also known as the Class of Service COS) and Type of Service (TOS). You can tell PacketShaper to match on IP precedence and/or TOS during classification of IP protocols, then apply specific policies to manage these traffic types. For details on using a policy to substitute an IP precedence value, see Substitute Diffserv Values.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To set up a Diffserv COS/TOS matching rule:

1. When creating a class or adding a matching rule, click **diffserv**. The *Matching Rule: Diffserv* window opens.

2. Select **COS/TOS** for the Diffserv Type. **show screen...**

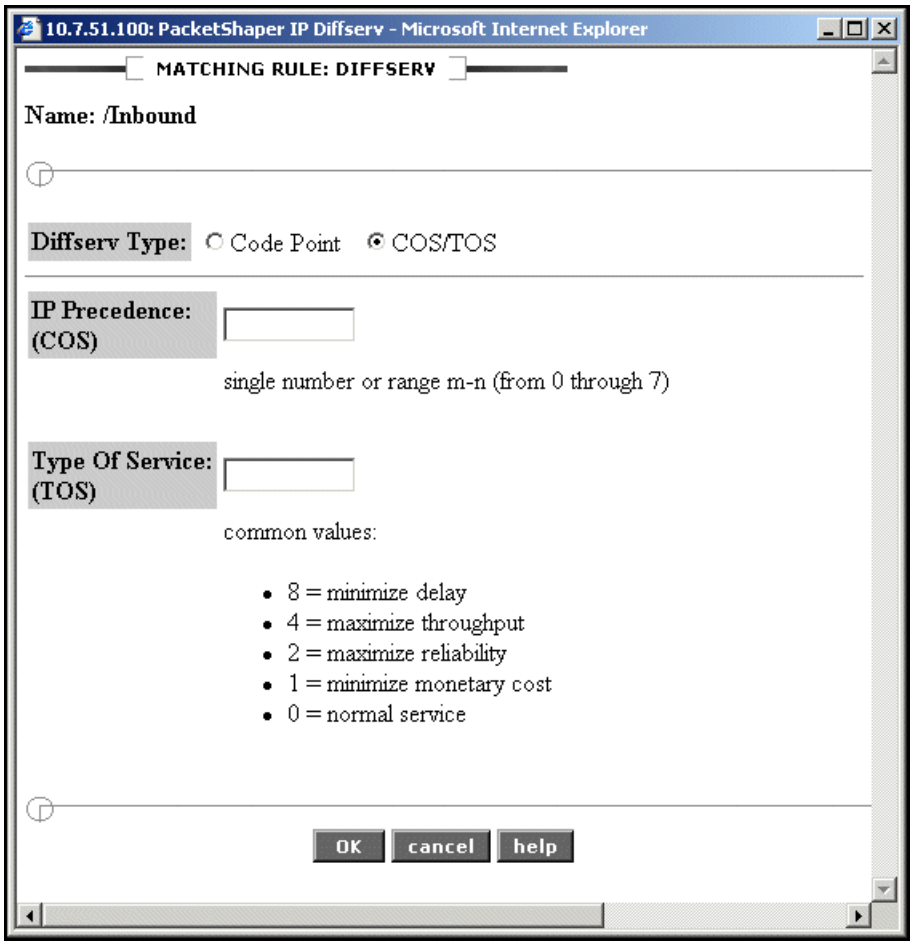

3. Enter the **IP Precedence (COS)** value (0 7) or a range of values.

4. Enter the **Type of Service (TOS)** value. Common values are listed on the screen.

5. Click **OK**.

6. Click **add class** or **add rule**.

Classify ICMP Messages

ICMP (Internet Control Message Protocol) provides message packets to report errors and other information relevant to IP packet processing. PacketShaper can classify ICMP traffic based on the message type. For example, the UNIX ping command generates ICMP echo request packets and listens for the subsequent ICMP echo response packets. To identify ping traffic passing through the unit, you can create an ICMP class and specify **echo** as the message type.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a class based on a type of ICMP message:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder that will be the parent of the class you are about to define.

3. Click **class** and select **add**.

4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).

5. For the **Service**, select **ICMP**.

6. Scroll down to the bottom of the page and make sure **ICMP Type** appears in the **Criterion** field. **show screen...**

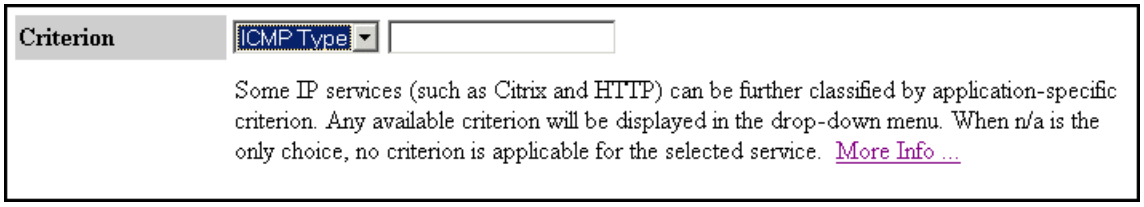

7. Type one of the following message type in the empty text field:

- echo
- dest_unreachable
- quench
- redirect
- time_exceeded
- param_prob
- timestamp
- info
- address_mask
- nachi

8. Click **add class**.

For more information about the other entries in the *New Traffic Class* window, see Matching Rules.

> **Note:** A matching rule can contain only one application-specific criterion. To combine multiple criteria, use two traffic classes, a parent and a child.

Classify RTP Traffic

If traffic discovery is enabled, PacketShaper will automatically discover Interactive Realtime Transport Protocol traffic, classify it in the RTP-I class, and create child classes for each discovered encoding name (such as G729 or JPEG) so that you can manage streams differently based on how they are encoded. Alternatively, if you want to auto-discover user agent child classes instead of encoding name classes, you can enable the enableVoIPUseragentAutoDiscovery system variable.

You can manually create RTP classes based on other attributes as well. You can classify by caller, callee, source IP address of the call setup (SIP) flow, or destination IP address of the flow. For example, suppose all VoIP calls go through a SIP Gateway as shown in the diagram below:

Host1 ------- SIP Gateway1 ---------- SIP Gateway2 ------- Host2

If you want to classify all VoIP traffic going through a specific SIP Gateway, you can create an RTP class with criteria that identifies the source or destination IP address of the SIP Gateway.

To create classes based on an RTP attribute, you must specify RTP-I as the service type.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To classify interactive RTP traffic:

1. Click the **Manage** tab.

2. In the left window pane, select the class or folder that will be the parent of the class you are about to define.

3. Click **class** and select **add**.

4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).

5. For the **Service**, select **RTP-I**.

6. Scroll down to the **Criterion** field, and select the desired criterion from the drop-down list and then enter a text string in the empty text field. For example, to create a class for a specific VoIP caller, select **SIP Caller Identifier** in the **Criterion** field and then type the identifier (such as *agb@bell-telephone.com*) in the text field.

> **Note:** For SIP attributes, you can enter a substring of the attribute. For example, to match all Motorola models, you can enter **Motorola** for the SIP User-Agent criteria.

Attribute	Examples of Values
Encoding Name	G729, GSM, JPEG, PCMA, PCMU **Note:** PCMA and PCMU are both specified in CCITT/ITU-T recommendation G.711. To specify G.711 encoding, enter either **PCMA** (Pulse Code Modulation a-law) or **PCMU** (Pulse Code Modulation mu-law) for the Encoding Name.

Attribute	Examples of Values
Media Type	**a** for audio, **v** for video
Clock Rate	Supported values include: 8000, 16000, 44100, 90000
SIP Callee Identifier	+12125551212@server.phone2net.com, 17476004249@10.10.254.56
SIP Caller Identifier	agb@bell-telephone.com, 17476004230@172.21.1.41
SIP User-Agent	*X*Lite*, *Asterisk*PBX*, *3CXVoipPhone** **Additional Information** ■ Use the class criteria recent CLI command to find the exact string to enter for the SIP User-Agent criterion. ■ If you want to auto-discover user agent child classes instead of encoding name classes, you can enable the enableVoIPUseragentAutoDiscovery system variable.
SIP srcIP Address	207.78.98.18
SIP destIP Address	207.78.98.18

8. Click **add class**.

For more information about the other entries in the *New Traffic Class* window, see Matching Rules.

> **Note:** A matching rule can contain only one application-specific criterion. To combine multiple criteria, use two traffic classes, a parent and a child.

Service Group Tasks

A *service group* is a set of services that have been grouped together based on a common functionality. For example, all services that classify peer-to-peer file sharing applications are in the P2P group. Service groups offer a quick way to classify and control a group of related services.

Create a Custom Service Group

If the built-in groups don't suit your needs, you can create your own groups. For example, if you have created user-defined services for your custom applications, you may want to create a custom group for them. You can create up to 25 custom service groups.

PolicyCenter

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a custom service group:

1. Click the **Setup** tab.

2. From the setup list, select **service groups**.

3. Make sure the **Groups** tab is selected. **show screen...**

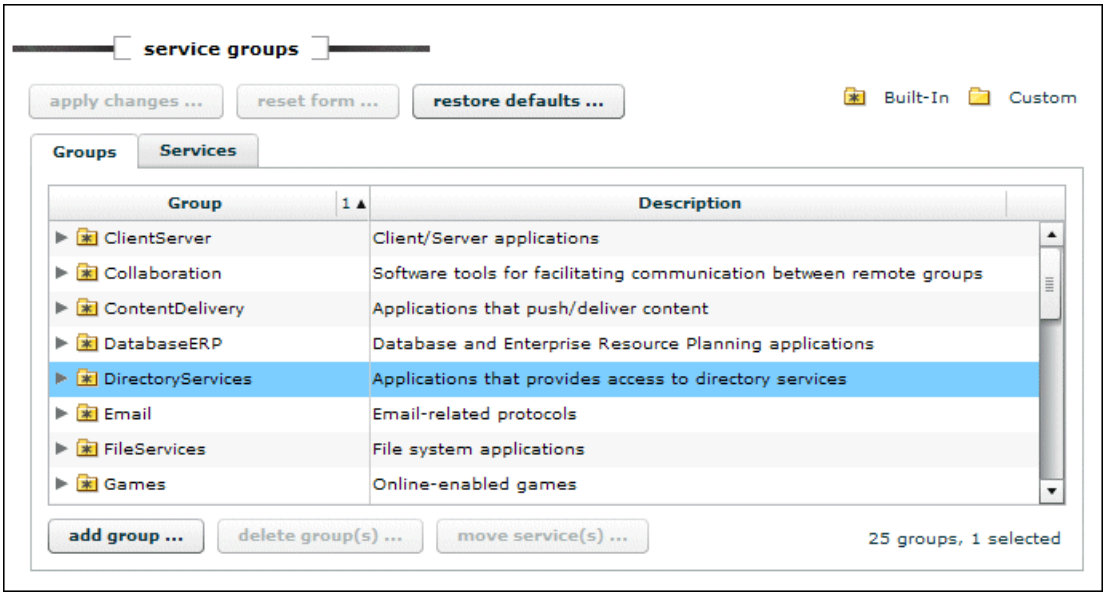

4. Click **add group**. The Add Group dialog box appears. **show screen...**

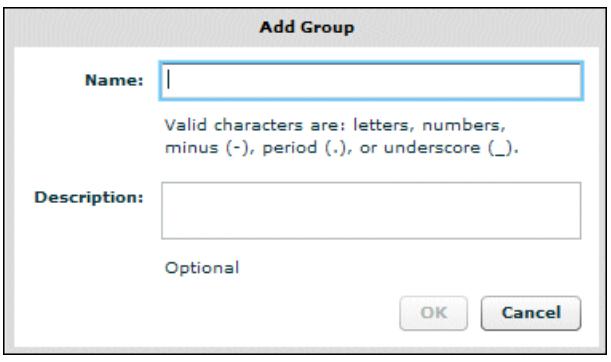

5. In the **Name** field, enter a descriptive name for the group (up to 31 characters; hyphens, underscores, and periods are acceptable).

6. In the **Description** field, enter text (up to 80 characters) to provide further information about the custom group.

7. Click **OK**. The new custom group appears at the bottom of the list; you may need to scroll down to see it.

8. Click **apply changes** to save the group.

> **Note:** After creating a custom group, you will probably want to move services into it.

Move Services to a Different Group

Although Symantec Systems has pre-defined a group for each service, you are free to move services into different built-in groups or into a custom group you have created. You can perform the move operation in either the Groups or Services tab.

PolicyCenter

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To move one or more services into a different group:

1. Click the **Setup** tab.

2. From the **Choose setup page** list, select **service groups**. **show screen...**

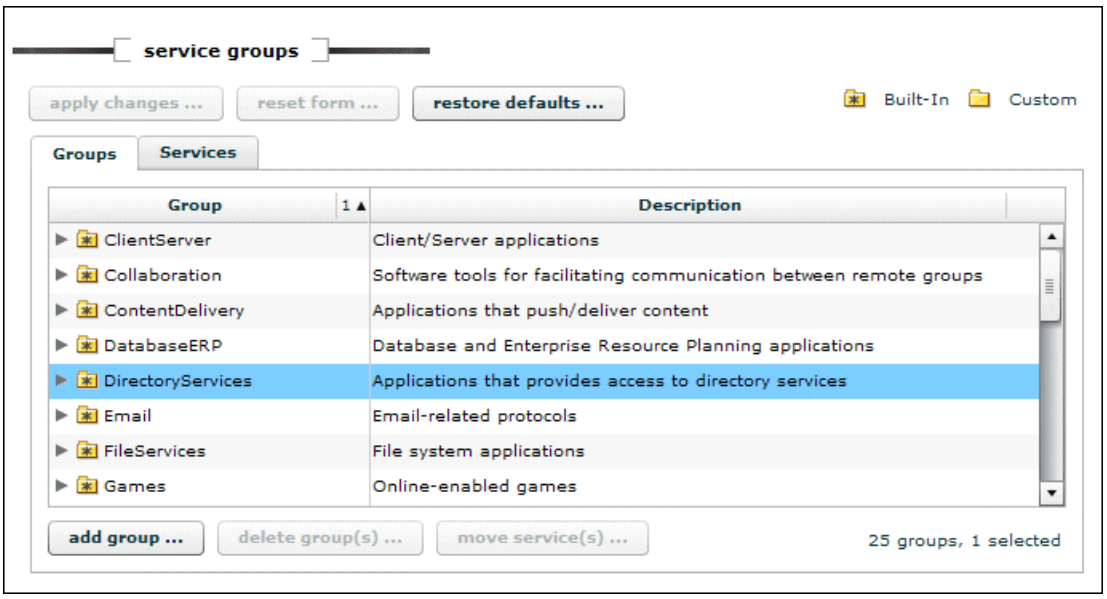

3. Decide whether you want to move services in the **Groups** tab or **Services** tab.

- If you know the name of the service you want to move but not the name of the group it currently belongs to, it's easiest to perform the move operation in the **Services** tab.
- If you want to move all the services in a particular group into another group, you will find it easier to select the services in the **Groups** tab.

4. Select the service(s) you want to move. To select multiple services, hold down Ctrl as you click each service name. If you are moving services in the Groups tab, you will need to open the group to see the services in the group (click the ▶ to the left of the group name).

> **Note:** The status line at the lower-right corner of the screen indicates the total number of services that are currently selected.

5. Click **move services**. The Move Services dialog box appears. **show screen...**

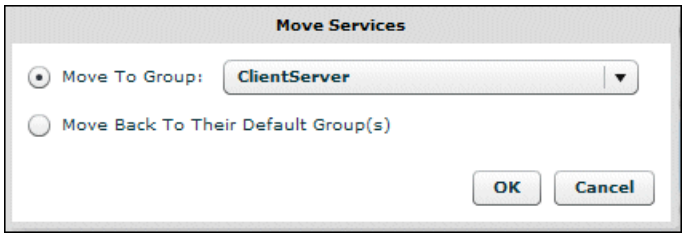

6. From the **Move to Group** list, select the name of the service group into which you want to move the selected services.

7. Click **OK**. The services are moved into the group you selected.

8. Click **apply changes** to save the change.

Additional Information

- If you later change your mind about the move, the Move Services dialog box contains an option to move services back into their default groups. See Returning Services to Default Groups.
- Services in the NonIPv4 group cannot be moved to other groups.

Return Services to Default Groups

The service group feature gives you the freedom to create custom groups as well as move services into different groups. If you change your mind, make a mistake, or just decide you want to go back to the factory-default settings, there are ways to recover.

Restore Settings to Factory Defaults

The **restore defaults** button on the Service Groups setup page moves all services back into their default built-in groups. If you had created any custom groups, these groups will remain, although they will no longer contain any services.

Use this button if you want to restore service groups to the factory defaults. If you want to return some services (not all) to their default groups, use the technique described below instead of the **restore defaults** button.

Move Services into their Default Groups

You don't need to know which group a service originally belonged to – you can just indicate that you want to return it to its default group.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To move one or more services to their default groups:

1. Click the **Setup** tab.

2. From the **Choose setup page** list, select **service groups**. **show screen...**

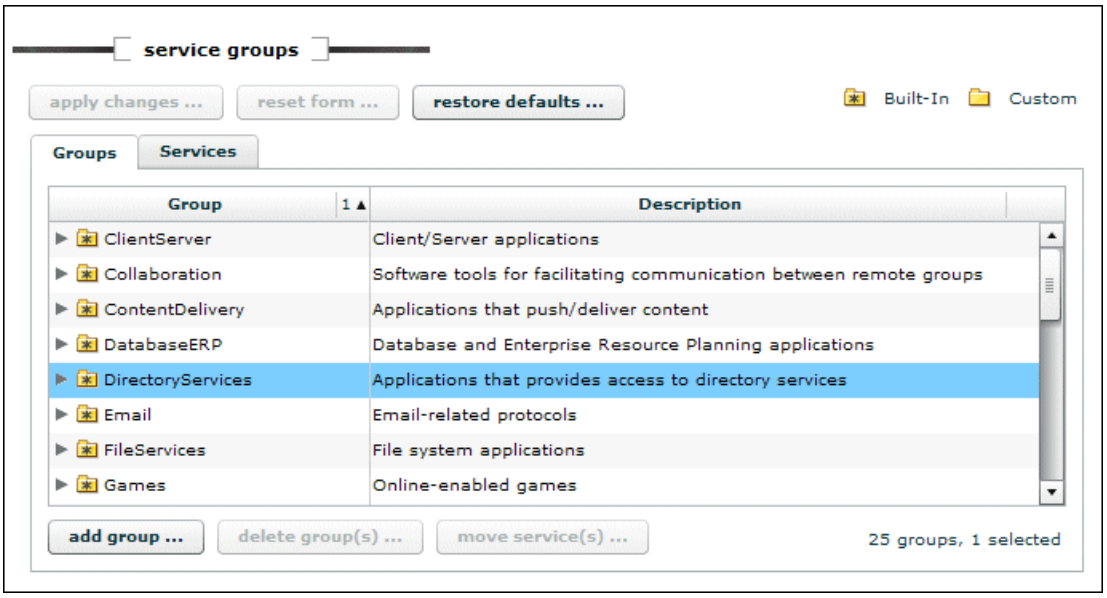

3. Decide whether you want to use the **Groups** tab or **Services** tab to move the services.

- If you know the name of the service you want to move, it's easiest to perform the move operation in the **Services** tab.
- If you want to return all the services in a particular group to their default groups, you will find it easier to use the **Groups** tab.

4. Select the service(s) you want to move. To select multiple services, hold down Ctrl as you click each service name. If you are selecting services in the Groups tab, you will need to open the group to see the services in the group (click the ▶ to the left of the group name).

5. Click **move services**. The Move Services dialog box appears. **show screen...**

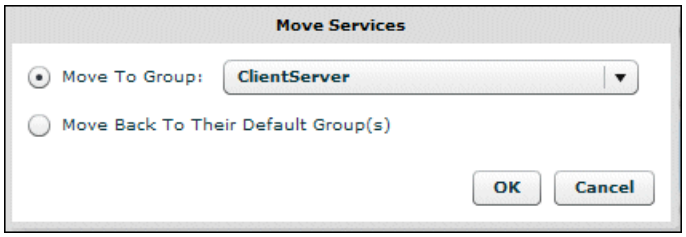

6. Select **Move Back To Their Default Group(s)**.

7. Click **OK**. The services are returned to their default groups.

8. Click **apply changes** to save your changes.

Delete a Custom Service Group

You can delete custom service groups, but you cannot delete the built-in groups. When you delete a group, the services in that group are moved into the Unassigned group.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To delete one or more custom service groups:

1. Click the **Setup** tab.

2. From the **Choose setup page** list, select **service groups**.

3. Make sure the **Groups** tab is selected. **show screen...**

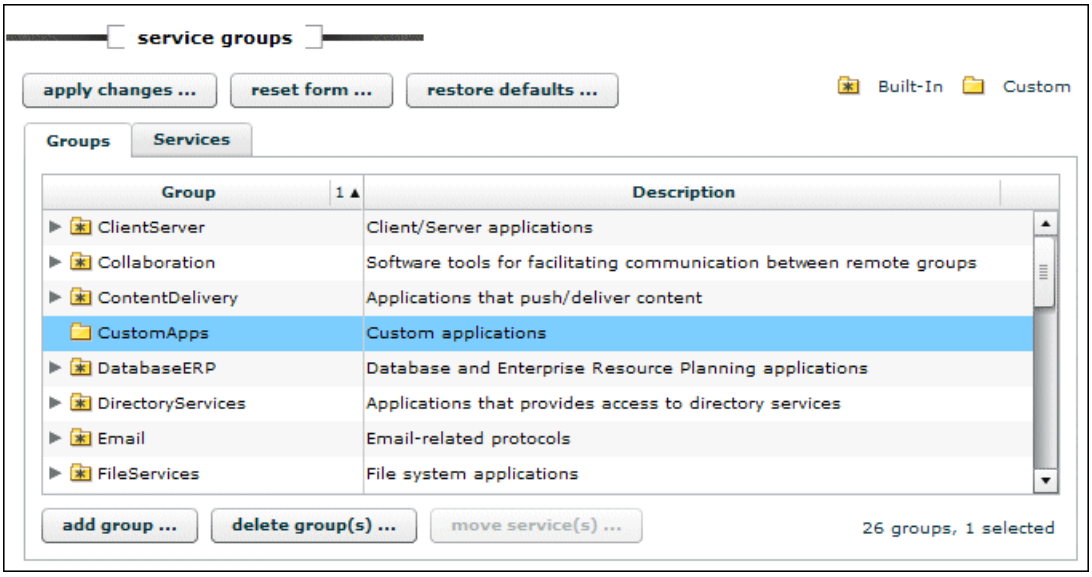

4. Select the custom group(s) you want to delete.

5. Click **delete group**. The Confirm Group Deletion dialog box appears.

6. Click **Yes**.

7. Click **apply changes** to save your change.

> **Note:** The **delete group** button is grayed out when no groups are selected or if built-in groups are selected.

Create a Service Group Class

In order to monitor and control traffic in a service group, you must create a traffic class for the group. You can create the service group class in either interface:

Create a Service Group Class in the Advanced UI

Create a Service Group Class in the Sky UI

Create a Service Group Class in the Advanced UI

In order to monitor and control traffic in a service group, you must create a traffic class for the group.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a service group class in the Advanced UI:

1. Click the **Manage** tab.

2. In the left window pane, select the class (such as Inbound or Outbound) that will be the parent of the class you are about to define.

3. Click **class** and select **add**.

4. In the **Name** field, type a descriptive name (up to 31 characters; hyphens, underscores, and periods are acceptable).

5. Select **Service Group** and choose the group name from the list. **show screen...**

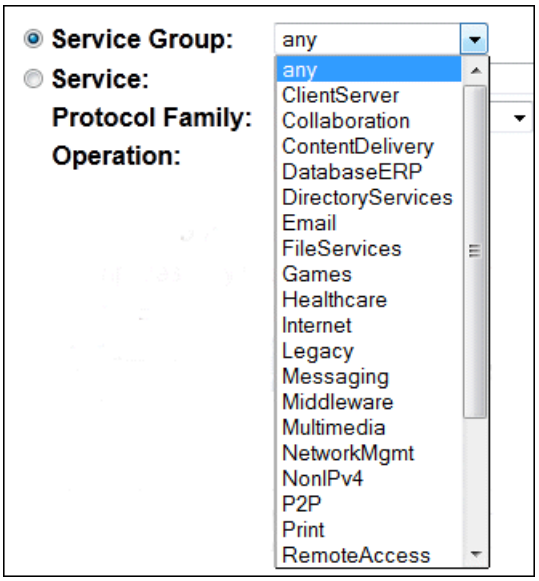

6. Click **add class**.

> **Note:** If you create a service group class (such as P2P) when the class tree already contains classes for individual services (such as eDonkey and Pando) in that group, you will need to delete the service-based classes so that new flows get classified into the service group class.

Create a Service Group Class in the Sky UI

In order to monitor and control traffic in a service group, you must create a traffic class for the group.

To create a service group class:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane.

2. In the class tree, select the parent class for your new class (such as Root).

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Add Class**.

6. In the **Name** field, type a descriptive name for the traffic class (up to 31 characters; hyphens, underscores, and periods are acceptable).

7. Select the **Direction** for the new class: **Both**, **Inbound**, or **Outbound**

8. If you want PacketShaper to discover subclasses within the class, enable the **Auto-Discovery in Class** checkbox.

9. From the **Type** drop-down list, select **Service Group**. **show screen...**

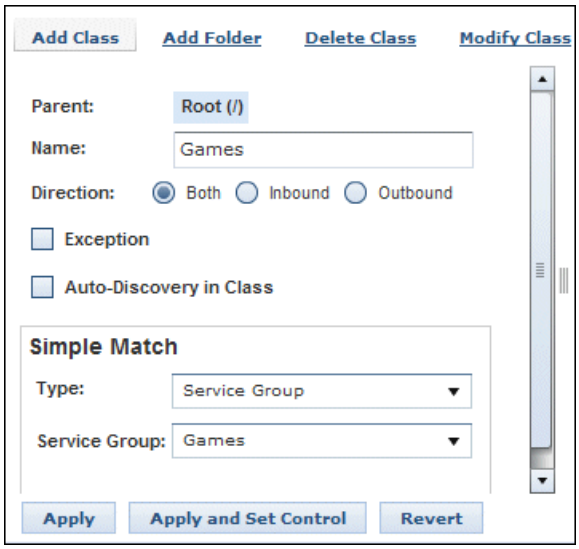

10. Choose the group name from the **Service Group** list.

11. Click **Apply** or **Apply and Set Control**.

> **Note:** If you create a service group class (such as P2P) when the class tree already contains classes for individual services (such as eDonkey and Pando) in that group, you will need to delete the service-based classes so that new flows get classified into the service group class.

Traffic Tree Management Tasks

To help you maintain and organize your traffic tree, PacketShaper offers tools for grouping classes into folders, moving classes to different branches in the tree, positioning classes closer to the top of the tree, and deleting and copying classes.

Collapse and Expand the Traffic Tree

The **manage** and **monitor** tabs display a traffic tree on the left side of the window. PacketShaper offers you a few ways to control which classes display on the traffic tree. If the tree has many classes, you may want to collapse parts of the tree so that you can focus on a single branch of the tree (such as Inbound, Outbound, or a particular folder).

Collapse the Tree

To collapse the Inbound and Outbound branches of the tree, click the **Collapse All** icon.

To collapse a branch of the tree, use one of the following techniques:

- Click the **Collapse Inbound** or **Collapse Outbound** icons.
- Click the **-** to the left of the folder you want to collapse.

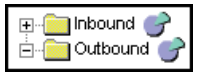

Expand the Tree

To expand the entire traffic tree, click the **Expand All** icon.

To expand a branch of the tree, use one of the following techniques:

- Click the **Expand Inbound** or **Expand Outbound** icons.
- Click the **+** to the left of the folder you want to expand.

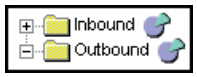

Set the Tree Depth

The **Tree Depth** field allows you to specify the number of class levels to be displayed in the traffic tree. For instance, with a tree depth of 1, the tree will display Inbound and Outbound and the first level of classes and folders under Inbound and Outbound. To set the level, select the number (1-8) from the **Tree Depth** drop-down list. The **off** setting expands the entire tree. **show screen...**

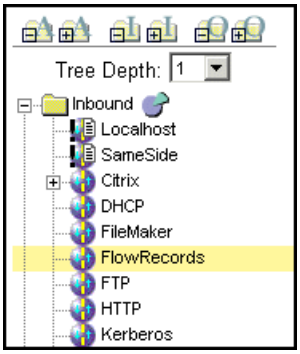

Regardless of the tree depth you have set, you can still manually collapse and expand individual branches of the tree.

Note that the **Tree Depth** value is retained within and between user sessions.

Large Traffic Trees

If you have more than 2000 classes in your configuration, PacketShaper does not display a traffic tree in the **Manage** tab. Instead of the tree, the left window pane displays a scrollable list of all the classes on your unit. To narrow down this list:

1. Type the name of the link, folder, or class in the field **Display only class names containing**.
2. Click **limit view**. The list now displays only the classes that contain the name you entered.

To display the entire list after limiting the view, click **display all**.

Create a Folder Class

You can use a folder class to organize the traffic tree. For example, you can create a folder called Critical_Apps and move all classes associated with critical applications into this folder. This structure allows you to see all your critical applications in one location in the traffic tree, plus it enables you to to create one partition to be used to protect critical traffic from other non-critical traffic. Two folders are created automatically: /Inbound and /Outbound. In addition, if port-based traffic is autodiscovered, the classes for this traffic are placed in DiscoveredPorts folders at the bottom of the /Inbound and /Outbound sections of the traffic tree.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a folder class:

1. Click the **Manage** tab. The *Traffic Class* window displays.

2. In the left window pane, select a class that will serve as the parent of the new folder class.

3. Click **class** and select **add folder**. **show screen...**

4. Specify the folder name and click **OK**.

Folder classes do not match any traffic that is, there are no specific traffic criteria associated with this type of class. Therefore, a policy cannot be applied to a folder class. You can, however, apply a partition to a folder.

Create an Exception Class

When PacketShaper is classifying a traffic flow, it searches the traffic tree, beginning at the top, to find a match for the flow. As soon as PacketShaper finds a match, it categorizes the flow into the appropriate class and looks no further down the tree (even if there may be another class that matches the traffic flow). When you want to change the order that PacketShaper uses to find a match for a traffic flow, you can create exception classes. An exception class sorts above standard classes in a traffic subtree. Exception classes give you the ability to redefine the search order that PacketShaper uses to find a match for a traffic flow. Typically, you do not need to change the position that PacketShaper defines for a class this is for special exceptions only (hence, the name *exception class*).

Certain classes (/Inbound, /Outbound, and Default classes) cannot be configured as exception classes.

> **Note:** Configuring too many classes as exceptions will degrade PacketShaper performance.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To designate an existing class as an exception:

1. Click the **Manage** tab.

2. In the left window pane, select the class you want to designate as an exception. The *Traffic Class* window displays options for the selected class. The default class type is Standard.

3. Next to Type, select **Exception**.

4. Click **apply changes**.

The class now appears at, or near, the top of the subtree. The "!" on the class icon indicates the class is an exception.

Create an Inheritable Traffic Class

Inheritable traffic classes have policies that can be applied to other classes when the other class doesn't have its own policy. Specific rules apply to how PacketShaper decides which policy a class should inherit; see Policy Inheritance Rules for details. The /Inbound/Default and /Outbound/Default classes are predefined as inheritable.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To designate an existing class as inheritable:

1. Click the **Manage** tab.

2. In the left window pane, select the class you want to designate as inheritable. The *Traffic Class* window displays options for the selected class. The default class type is Standard.

3. Next to Type, select **Inheritable**.

> **Note:** If the class does not list Inheritable as an option, it doesn't have a policy. Only classes with a policy can be designated as inheritable.

4. Click **apply changes**.

The class name now appears bolded in the traffic tree, indicating that it has an inheritable policy.

Inherited Class Example

ISPs can create classes for certain applications (such as Inbound/HTTP), assign appropriate policies, and designate these classes as inheritable. In this scenario, traffic is first classified by customer IP address or subnet and allocated bandwidth (via a partition). Then, the traffic is assigned a policy according to the type of flow. For example, if the inbound traffic is HTTP, it will be assigned the Inbound/HTTP policy (since this class was designated as inheritable and is a sibling class of the customer partition). With this configuration, the ISP can not only ensure each of its customers is getting its share of bandwidth, but can also improve application performance for its customers. Here s what this type of traffic tree might look like:

Inbound

- Customer1 (partition)
- Customer2 (partition)
- Customer3 (partition)
- Customer4 (partition)
- **HTTP** (inheritable class with a policy; matches all HTTP flows)
- **Default** (inheritable class with a Priority 3 policy, matches all non-HTTP flows)

Update Class Properties

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To display or update the properties of a traffic class:

1. In the **Manage** tab, select a class in the left window pane.

When a class is selected, its details are displayed in the *Traffic Class* window. A menu with four options appears at the top: **class**, **policy**, **partition**, and **statistics**. **show screen...**

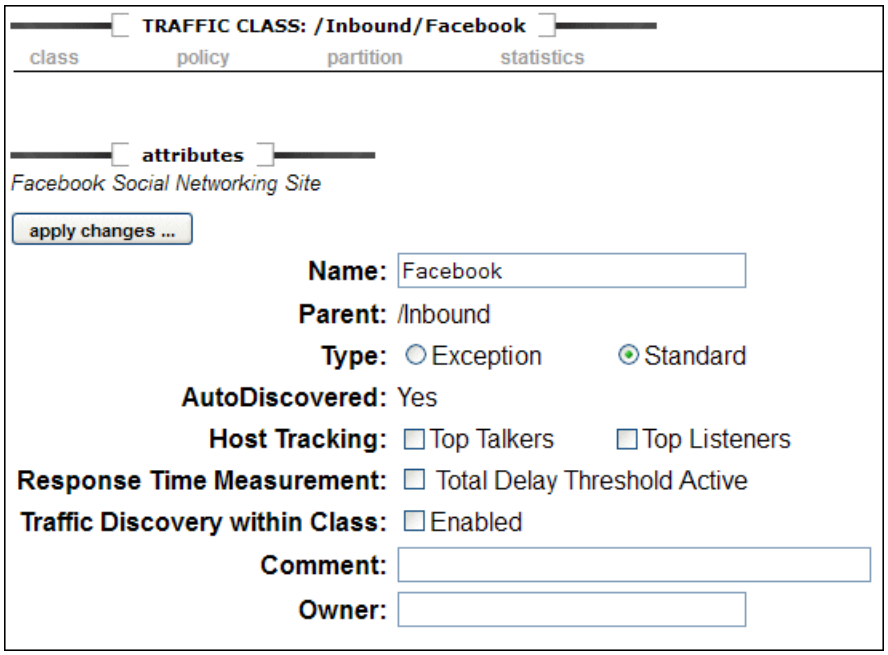

2. Verify or modify properties, as described in the following table.

3. Click **apply changes** to update the settings.

Class Property	Description
Name	Rename a traffic class by giving it a unique name (up to 31 characters). Use only alphanumeric characters and underscores (_), dashes (-), and periods (.). Note that when renaming a class you are not allowed to change just the case; for example, you cannot rename HTTP to http.
Parent	The class' parent in the traffic tree is displayed (cannot be modified).

Class Property	Description
Type	**Exception** traffic classes are always positioned above non-exception classes in the tree. When you make a class an exception class, you redefine the search order that PacketShaper uses to find a match for traffic flow. For details, see Create an Exception Class. **Standard** traffic types have no exception or inheritable attributes. **Inheritable** traffic classes have policies that can be applied to other classes when the other class doesn't have its own policy. For details, see Create an Inheritable Traffic Class.
AutoDiscovered	If **Yes** is displayed here, the traffic class was created through PacketShaper's automatic traffic discovery feature. If **No** is displayed, the class was manually created. This field cannot be modified.
Traffic Discovery within Class	Enable/disable traffic discovery within a specific class. Some traffic classes will not allow traffic discovery within the class because they already are as detailed as possible—that is, child classes cannot be generated for the class. For example, a class that matches on IPX cannot be further classified. For details, see Turn on Traffic Discovery for an Individual Class.
Host Analysis	Check the **Top Talkers** and **Top Listeners** checkboxes to track the top twenty talkers and/or listeners for this class. When these options are selected, they become links that take you to other windows. For details, see Track Hosts that Generate the Most Traffic.
Comment	Enter a description (up to 79 characters). **Note:** The following special characters are not allowed: double quote ("), single quote ('), ampersand (&), backslash (\), vertical bar (|).
Owner	If this class is associated with a particular user or customer, specify the customer name (up to 32 characters). **Note:** The following special characters are not allowed: double quote ("), single quote ('), ampersand (&), backslash (\), vertical bar (|).
Performance	This item appears when the selected address-based class is not currently cached. (The name of an uncached class also appears in italic in the traffic tree.) Because uncached classes can cause a minor degradation in performance, you can improve performance by making all address-based classes cacheable. One way to do this is to remove a match-all class near the uncached class. For example, try removing a Default bucket that is a sibling of the parent of the uncached class(es).
Matching Rule	A class' matching rule displays the characteristics of the selected traffic class. See Matching Rules.

Test a Traffic Class

To check a particular IPv4 or IPv6 address and/or port number against the classification tree to determine how the traffic is classified, you can go to the *Test Traffic Class* window.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To test how traffic will be classified:

1. Click the **Manage** tab.

2. Click the **traffic class test** button. The *Test Traffic Class* window appears. **show screen...**

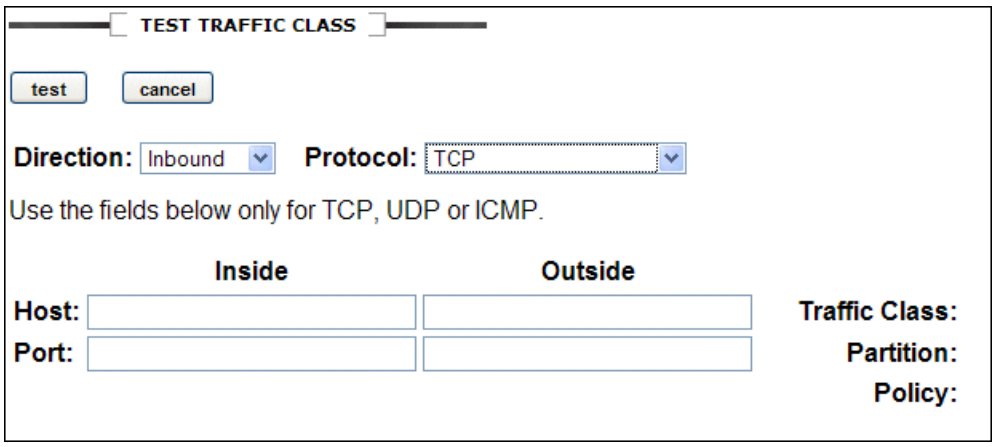

3. In the *Test Traffic Class* window, in the **Direction** field, select **Inbound** or **Outbound** to determine which portion of the classification tree is to be tested.

4. Select the protocol: TCP, UDP, ICMP, IPX, SNA, AppleTalk, NetBEUI, DECNET, FNA, Others (Misc).

5. **IPv4 hosts:** For TCP, UDP, and ICMP protocols, specify the Inside and Outside IPv4 address and/or port number of the hosts. You must supply both an inside and an outside address. Use 0.0.0.0:0 as a placeholder if you don't have an address to test on one of the sides.
or
IPv6 hosts: For TCP and UDP protocols, specify the Inside and Outside IPv6 address and/or port number of the hosts. You must supply both an inside and an outside address.

6. Click **test**.

PacketShaper displays how the address was matched against the traffic tree. The results include the traffic class name, partition name, and the policy. **show screen...**

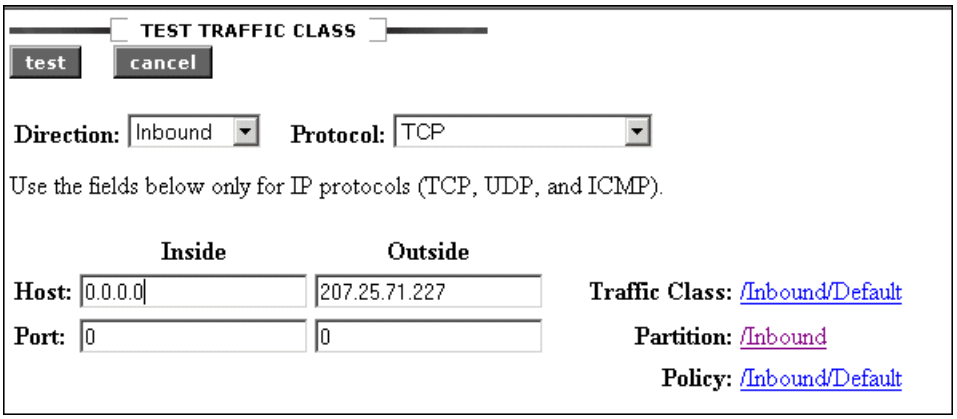

If the traffic class does not have a policy directly associated with it, it inherits a policy from an inheritable sibling class. (See Inheritance Rules.)

Additional Information

- The Traffic Class Test command can be used to test basic classification for IP protocols, but is not intended to test every type of classification PacketShaper offers. Its purpose is to check a particular IP address or port number to determine how the traffic is classified into existing port-based and IP address-based classes in the traffic tree. The command does not include fields for specifying more complex types of classification such as MAC address, although the class test CLI command has an option for specifying a device name.

- The Traffic Class Test command will only match traffic classes that have "any" for the server location.

Copy a Class

When you need to create a class that has similar specifications and matching rules as an existing class, you can copy a class. You can copy a class to a different subtree only; you cannot copy a class to the same parent.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To copy a class to a different subtree:

1. Click the **Manage** tab.

2. In the left window pane, select the class you want to copy.

3. In the right window pane, click **class** and select **copy**. The *Copy Traffic Class* screen appears. **show screen...**

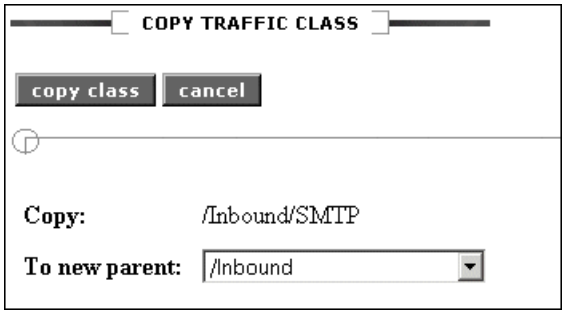

4. In the **To new parent** drop-down list, select the destination parent class.

5. Click **copy class** to copy the class under the designated new parent.

or

Click **copy class and children** to copy the class and its children to the designated parent.

To copy multiple classes at one time, see Copy Multiple Classes.

> **Note:** The **copy class** option does not copy the Traffic Discovery within Class attribute; the **copy class and children** option does copy the attribute.

Copy Multiple Classes

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To copy multiple classes:

1. Click the **Manage** tab.

2. Click **class** and select **bulk operation**. The *Traffic Class Bulk Operation* window appears. **show screen...**

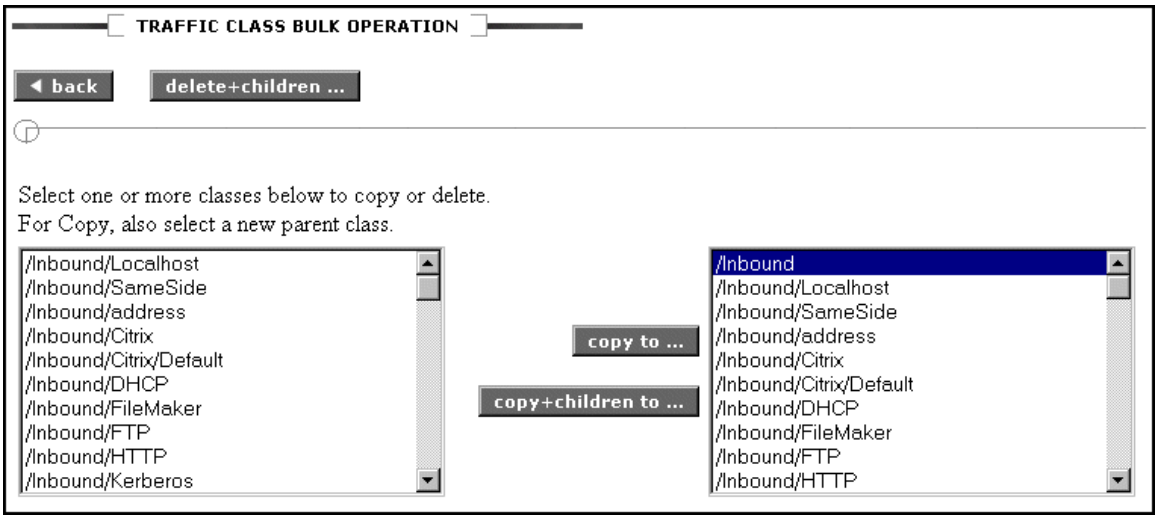

3. Select multiple classes from the left-hand list using one of the following selection methods:

   - Ctrl+Click to select noncontiguous classes
   - Shift+Click to select a set of contiguous classes

4. Select a single destination parent class from the list at the right of the window.

5. Click **copy to** (to copy only the classes you selected) or **copy+children to** (to copy the selected classes and their child classes).

Delete a Class

When you delete a class, the class name is removed from the traffic tree and PacketShaper will no longer classify flows for the associated service or application. Keep in mind, though, that if traffic discovery is enabled, the discovery process may recreate the class based on the traffic it discovers.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To delete a traffic class:

1. Click the **Manage** tab.

2. In the left window pane, select the class you want to delete.

3. In the right window pane, click **class** and select **delete** or **delete+children**.

> **Note:** If the class has children, you must choose the **delete+children** option. You cannot delete a class without deleting its child classes.

To delete multiple classes at one time, see Delete Multiple Classes.

Delete Multiple Classes

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To select multiple classes for deletion:

1. Click the **Manage** tab.

2. Click **class** and select **bulk operation**. The *Traffic Class Bulk Operation* window appears. **show screen...**

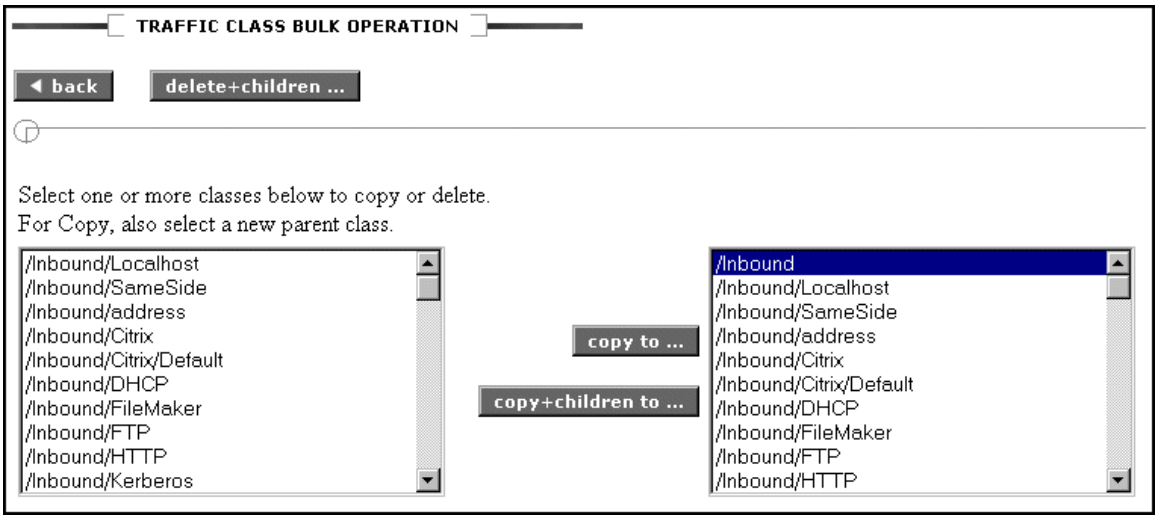

3. Select multiple classes from the left-hand list using one of the following selection methods:

   - Ctrl+Click to select noncontiguous classes
   - Shift+Click to select a set of contiguous classes

4. Click **delete+children**.

The classes you selected, as well as any child classes (whether or not you actually selected the child classes) will be deleted.

Move a Class

When you want to reorganize the traffic tree, you can move a traffic class to a different parent traffic class or folder.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To move a class to a different subtree:

1. Click the **Manage** tab.

2. In the left window pane, select the class you want to move.

3. In the right window pane, click **class** and select **move**. The *Move Traffic Class* screen appears. **show screen...**

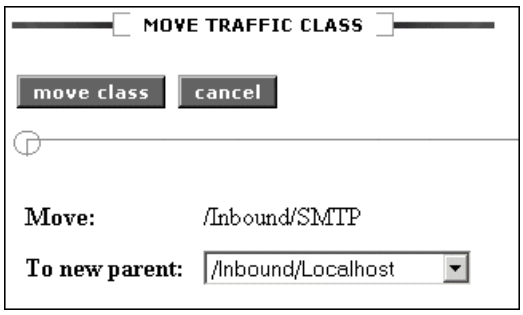

4. In the **To new parent** drop-down list, select the destination parent class.

5. Click **move class** to move the class under the designated new parent.

PacketShaper verifies that the moved class is reachable before completing the move. For example, if you attempt to move a TCP-based class to an IPX parent, the software notifies you that the move is not allowed.

When you move a class to a parent that has no children, a Default, match-all class is created automatically at the same time.

Moving a class moves its statistics along with it. If you are creating different policies and partitions for the moved class, you may want to clear the statistics since they may no longer be applicable.

Monitoring Tasks

Once PacketShaper has discovered and classified your network traffic, you can monitor the traffic to determine how bandwidth is used. You can find out which applications are consuming the most bandwidth and what percent of the link various applications are using. PacketShaper can quickly answer such questions as:

- What is the bandwidth usage of each mission-critical application?
- What portion of the link is consumed by non-critical or unsanctioned traffic?
- Who are the top network users?

Monitor Traffic in the Advanced UI

You can display the *Monitor Traffic* window in the Advanced UI by clicking the **Monitor** tab. **show screen...**

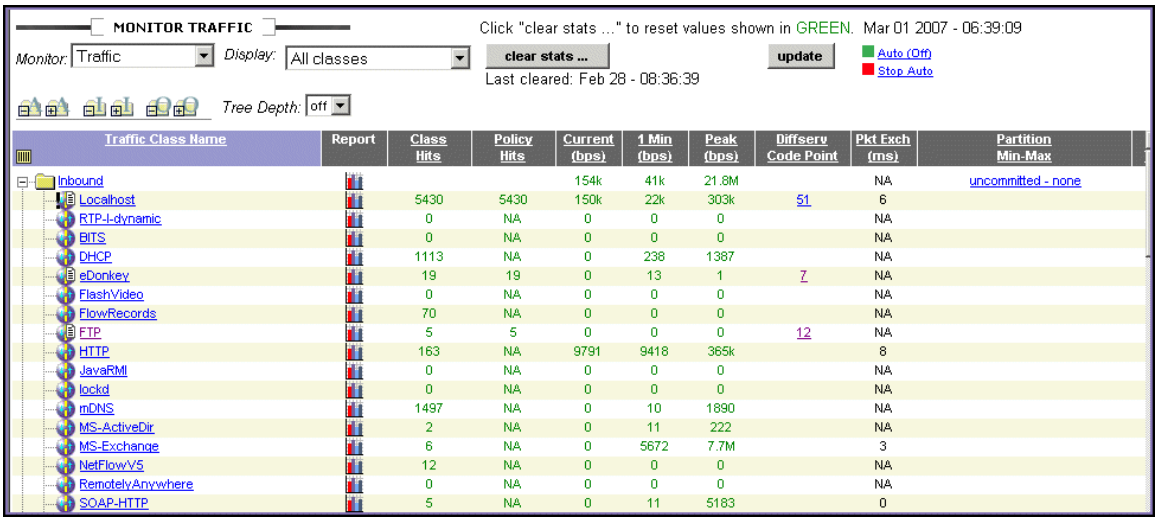

Traffic Class Name	Report	Class Hits	Policy Hits	Current (bps)	1 Min (bps)	Peak (bps)	Diffserv Code Point	Pkt Exch (ms)	Partition Min-Max
Inbound				154k	41k	21.8M		NA	uncommitted - none
Localhost		5430	5430	150k	22k	303k	51	6	
RTP-I-dynamic		0	NA	0	0	0		NA	
BITS		0	NA	0	0	0		NA	
DHCP		1113	NA	0	238	1387		NA	
eDonkey		19	19	0	13	1	7	NA	
FlashVideo		0	NA	0	0	0		NA	
FlowRecords		70	NA	0	0	0		NA	
FTP		5	5	0	0	0	12	NA	
HTTP		163	NA	9791	9418	365k		8	
JavaRMI		0	NA	0	0	0		NA	
lockd		0	NA	0	0	0		NA	
mDNS		1497	NA	0	10	1890		NA	
MS-ActiveDir		2	NA	0	11	222		NA	
MS-Exchange		6	NA	0	5672	7.7M		3	
NetFlowV5		12	NA	0	0	0		NA	
RemotelyAnywhere		0	NA	0	0	0		NA	
SOAP-HTTP		5	NA	0	11	5183		0	

The *Monitor Traffic* window lists the following information for each traffic class (unless you have chosen not to display certain columns):

Column	Description
Class Hits	Number of times flows match the class. Class hits occur only at the beginning of a flow or session. Classes created by traffic discovery do not show the first few connections for each traffic type. These initial hits are counted in the Default class. **Note:** In a configuration in which the PacketShaper unit sits between subnets and a router, traffic between two inside hosts will pass through the unit twice: once toward the router, then again back from the router to the other inside host. This results in double hit counts on a SameSide class. For example, four ping packets from one inside host to another would create eight class hits, as opposed to four with a standard setup (inside host to outside host).
Policy Hits	Number of times the policy has been enforced
Current (bps)	The current rate values in bits per second
1 Min (bps)	Last minute's average rate
Peak (bps)	The current peak values **Note:** Partitions have a 6% tolerance. Therefore, it is possible for an inbound traffic class to have a peak that exceeds the overall inbound peak.

Column	Description
Diffserv Code Point	The configured Diffserv policy for the traffic class. (Click the value to modify the policy.)
Pkt Exch (ms)	The average packet exchange time, which is the packet exchange time divided by the number of samples. NA means no packet exchange time samples were recorded. The *packet exchange time* is the interval between a data packet leaving the PacketShaper and the arrival of its acknowledgment (ACK). This variable is sampled every minute. Acknowledgments to synchronization flags (SYNs) are not included in this measurement.
Partition Min-Max	The minimum and maximum amount of bandwidth allocated for a class subtree. (Click the partition to modify it.) *Uncommitted* is the difference between the current link speed and the sum of all committed partition minimums. *Uncommitted* indicates that the partition size is dynamically determined at run time and the minimum allocation is whatever is not committed to other partitions. *None* indicates that the partition is burstable without a limit. If the min and max values are the same, the partition is fixed and non-burstable. If a percentage appears after the partition size, the partition was specified as a percentage of the parent partition's size. Example: *450k (10%)* for the Inbound/HTTP class indicates that HTTP's partition size is 10% of the Inbound partition's size, resulting in a partition size of 450 Kbps (assuming the link size is 4.5 Mbps). If a down arrow appears next to the partition size, the partition is dynamic. Place the mouse pointer on the down arrow and the parameters of the dynamic partition will be displayed: minimum and maximum subpartition size, the number of active and idle users currently using subpartitions, and the maximum number of subpartitions available.
Policy (Pri.) Guar.-Limit	Shows the policy (if any) applied to the class. (Click the policy to modify it.) For a priority policy, the defined priority is shown. For a rate policy, the guaranteed rate setting, cap, and burstable priority level are shown.
Top Host Tracking	Microphone and ear icons indicate if Top Talkers and Top Listeners are being tracked for the class. See List Top Hosts or Users . To view the top bandwidth talkers and listeners, click the respective icons to display a page of transmitter or receiver statistics.

Standby Feature

When the Standby feature is enabled, and a connection is set up with a partner, an additional tab appears in the Advanced UI. This tab is labeled with the IP address of the partner PacketShaper and contains the monitor tab of the partner. Use this tab to monitor traffic of the standby partner. As traffic goes through the partner unit, the statistics automatically update. See " Verify Standby Partners" on page 724 for more information. **show screen...**

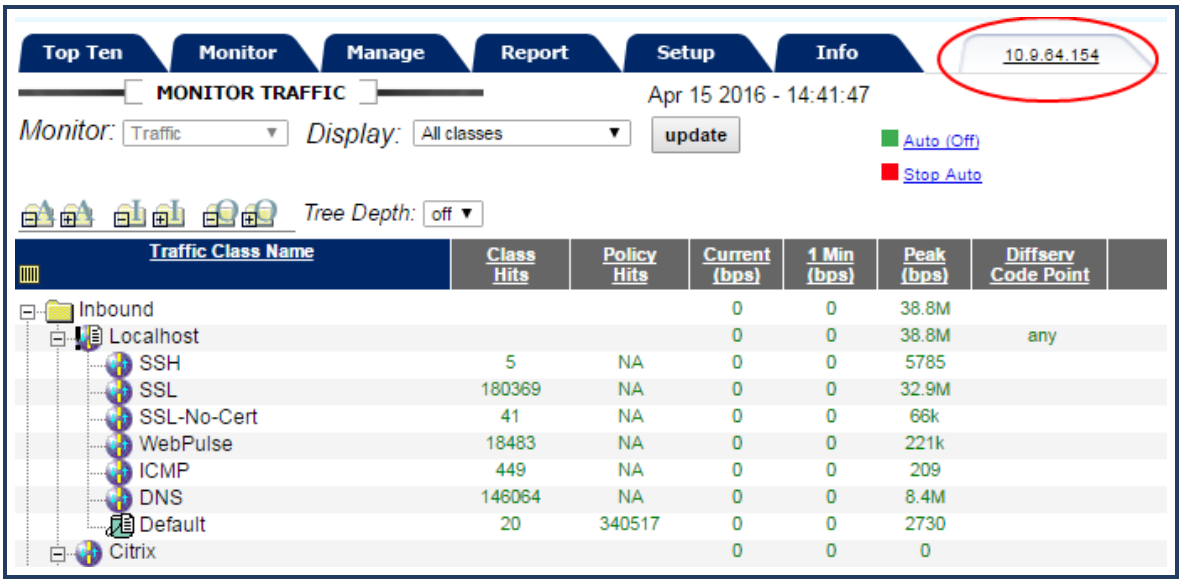

When a PacketShaper takes over management of a failed partner, the partner monitor tab (labeled with its IP address) for the failed unit disappears from the Advanced UI of the active PacketShaper(s) it has partnered with. During the failure, the Monitor tab in the active PacketShaper contains the traffic tree of the failed partner. The failed partner's inbound tree is at the bottom of the Inbound branch of the active PacketShaper, and the outbound tree is under the Outbound branch. The failed partner's tree is in a folder labeled with SB_ and the appliance serial number (for example SB_031400000). For more information, see "Standby Partner Failure" on page 726. **show screen...**

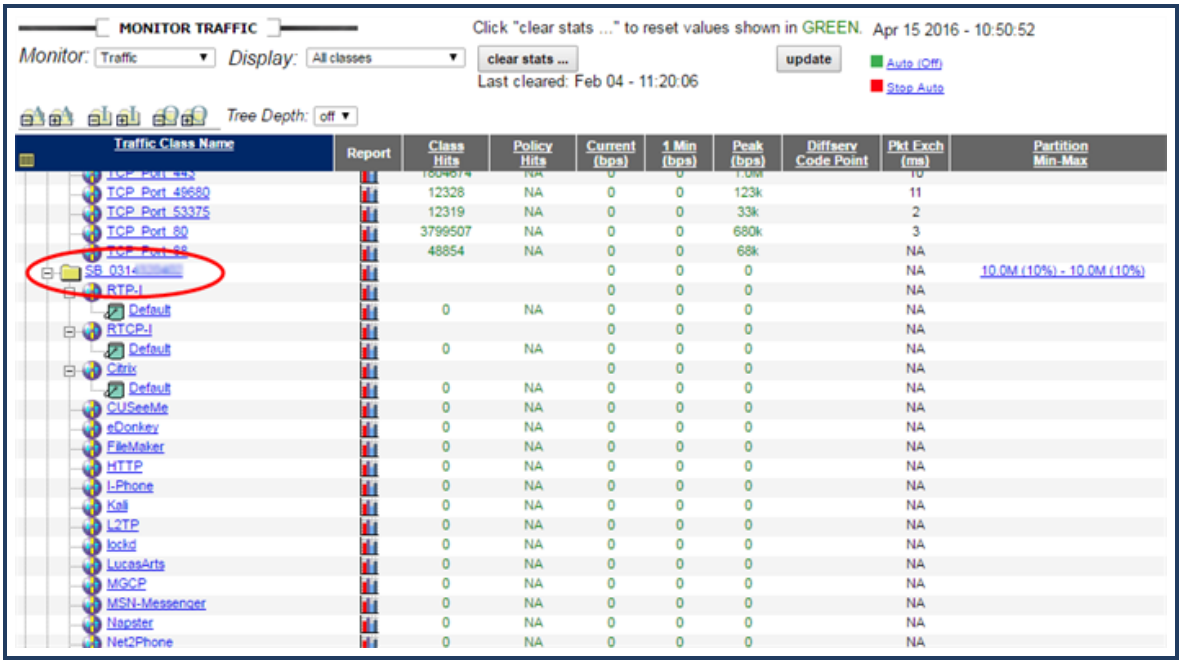

Select and Sort

The *Monitor Traffic* window, by default, lists all classes in your traffic tree, sorted by class name, and provides a variety of statistics (such as class hits, current rate, and peak rate) and settings (such as policies and partitions) associated with each of these classes. **show screen...**

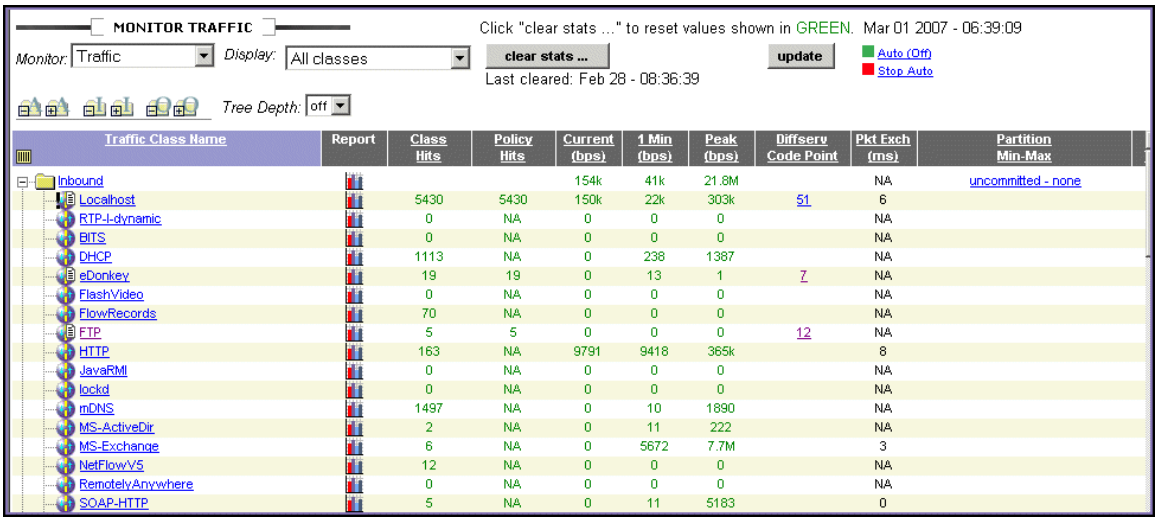

However, you have the ability to select which classes are displayed as well as which selectively list and sort columns. This capability gives you the power to zoom in on the information you need.

Selecting Columns to Display

You can customize the *Monitor Traffic* window to display only the columns of interest to you.

To select the columns to display in the *Monitor Traffic* window:

1. Click the **Monitor** tab. The *Monitor Traffic* window appears.

2. Click the **Select Columns** button (under the Traffic Class Name column heading). A list of column heading names appears; all columns are selected initially. **show screen...**

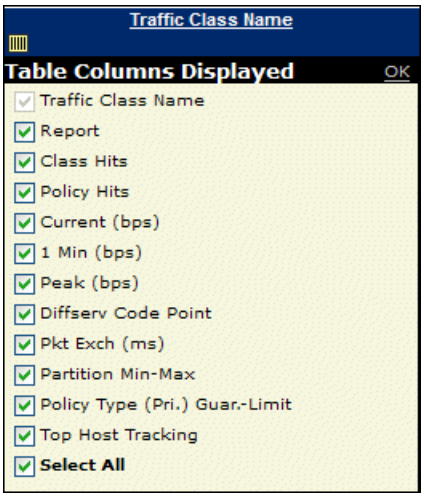

3. Select or clear checkboxes until the list contains the columns you want to display in the *Monitor Traffic* window.

or

If you want to display all columns, choose the **Select All** checkbox.

4. Click **OK**.

PacketShaper will retain the columns settings throughout your current login session only. The settings are not retained after you close the browser window.

Sorting the Monitor Traffic Table

The *Monitor Traffic* window is sorted by the Traffic Class Name column, but you can sort this table by any column. Just click the heading of the column you want to use for the sort. The information in the *Monitor Traffic* window is arranged from highest to lowest according to that column. For example, if you select the Class Hits heading, the summary is rearranged according to the number of hits a traffic class has incurred, with the most hits first.

Selecting Rows to Display

The *Monitor Traffic* window can be set to show only the classes that have partitions, policies, or are active. To change the list of classes displayed:

1. Click the **Monitor** tab. The *Monitor Traffic* window appears.

2. In the **Display** drop-down list, select one of the following:

- **All Classes**. Lists all classes in the traffic tree.
- **Active Classes**. Lists all classes into which traffic is currently being classified (that is, classes that are currently experiencing hits).
- **Classes with Partitions**. Lists all classes that have partitions assigned to them.
- **Classes with Policies**. Lists all classes that have policies assigned to them.

PacketShaper will retain the display setting – even if you go to another tab or exit your session.

Reset Monitor Statistics

The following real-time statistics are listed in the *Monitor Traffic* window: **show screen...**

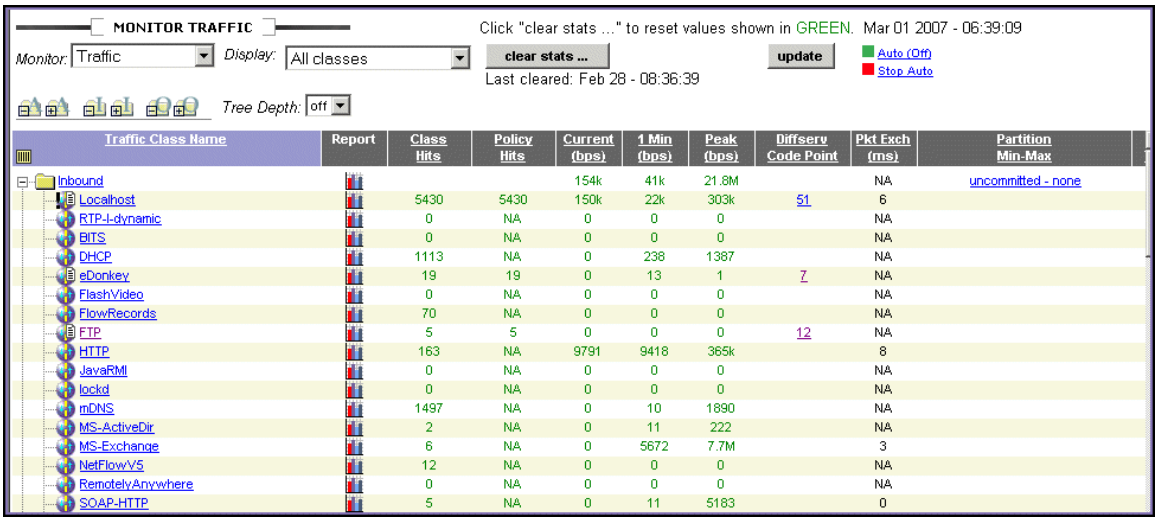

Traffic Class Name	Report	Class Hits	Policy Hits	Current (bps)	1 Min (bps)	Peak (bps)	Diffserv Code Point	Pkt Exch (ms)	Partition Min-Max
Inbound				154k	41k	21.8M		NA	uncommitted - none
Localhost		5430	5430	150k	22k	303k	51	6	
RTP-I-dynamic		0	NA	0	0	0		NA	
BITS		0	NA	0	0	0		NA	
DHCP		1113	NA	0	238	1387		NA	
eDonkey		19	19	0	13	1	7	NA	
FlashVideo		0	NA	0	0	0		NA	
FlowRecords		70	NA	0	0	0		NA	
FTP		5	5	0	0	0	12	NA	
HTTP		163	NA	9791	9418	365k		8	
JavaRMI		0	NA	0	0	0		NA	
lockd		0	NA	0	0	0		NA	
mDNS		1497	NA	0	10	1890		NA	
MS-ActiveDir		2	NA	0	11	222		NA	
MS-Exchange		6	NA	0	5672	7.7M		3	
NetFlowV5		12	NA	0	0	0		NA	
RemotelyAnywhere		0	NA	0	0	0		NA	
SOAP-HTTP		5	NA	0	11	5183		0	

- Class and Policy Hits
- Current, 1 Min, and Peak (bps)
- Guaranteed Rate Failures

The real-time statistics are shown in green type on the screen.

To reset the real-time statistics, click **clear stats**. All values on the screen are zeroed out. Occasionally, some statistics do not have zero values, even after you click **clear stats**; this happens because the unit may be logging traffic flows during the clearing of statistics.

Under the **clear stats** button, there is a line of text indicating the last date/time that the statistics were reset. This updates when the **clear stats** button is pressed. If the button has not been pressed, it shows the last time the PacketShaper was reset, since that also clears the statistics.

Update the Monitor Statistics

The statistics in the *Monitor Traffic* window are captured the moment you click the **Monitor** tab. To get a new snapshot, click **update** at the top of the screen. **show screen...**

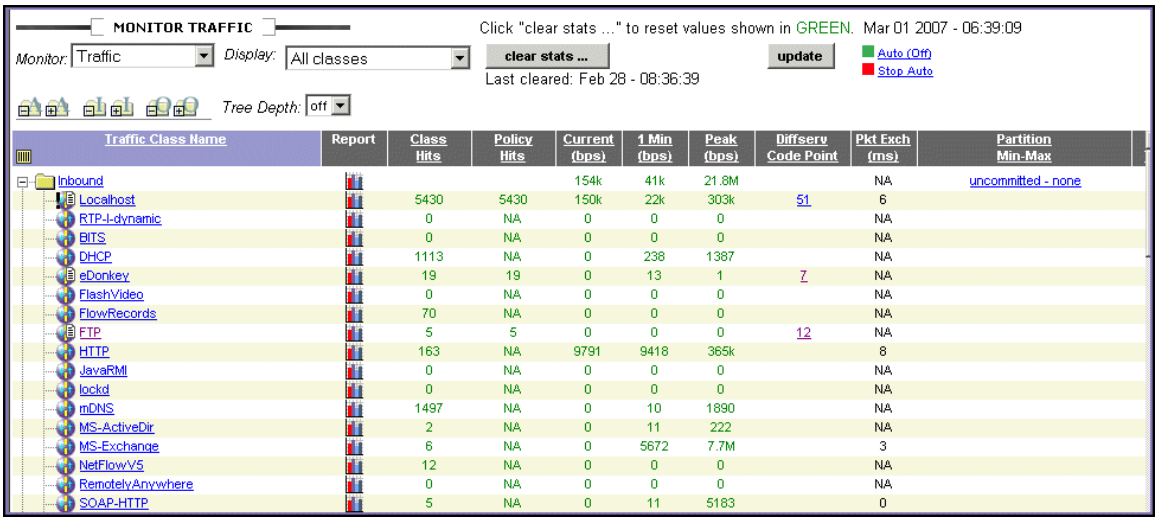

Traffic Class Name	Report	Class Hits	Policy Hits	Current (bps)	1 Min (bps)	Peak (bps)	Diffserv Code Point	Pkt Exch (ms)	Partition Min-Max
Inbound				154k	41k	21.8M		NA	uncommitted - none
Localhost		5430	5430	150k	22k	303k	51	6	
RTP-I-dynamic		0	NA	0	0	0		NA	
BITS		0	NA	0	0	0		NA	
DHCP		1113	NA	0	238	1387		NA	
eDonkey		19	19	0	13	1	7	NA	
FlashVideo		0	NA	0	0	0		NA	
FlowRecords		70	NA	0	0	0		NA	
FTP		5	5	0	0	0	12	NA	
HTTP		163	NA	9791	9418	365k		8	
JavaRMI		0	NA	0	0	0		NA	
lockd		0	NA	0	0	0		NA	
mDNS		1497	NA	0	10	1890		NA	
MS-ActiveDir		2	NA	0	11	222		NA	
MS-Exchange		6	NA	0	5672	7.7M		3	
NetFlowV5		12	NA	0	0	0		NA	
RemotelyAnywhere		0	NA	0	0	0		NA	
SOAP-HTTP		5	NA	0	11	5183		0	

You can set the *Monitor Traffic* window to automatically update at regular intervals. To enable auto-update:

1. Click **Auto**. You are prompted for an interval.

2. Enter the desired number of seconds between updates, and click **OK**.

To stop automatic updating, click **Stop Auto**.

To increase the update interval, do one of the following:

- Click **Auto**. Each time you click **Auto**, three seconds are added to the interval.
- Click **Stop Auto** to disable automatic updates, and then click **Auto** to enter a new interval.

Track Hosts and Users that Generate the Most Traffic

PacketShaper keeps track of the IPv4 and IPv6 hosts that generate the most traffic. You can configure PacketShaper to track the Top Talkers (hosts that initiate the most traffic) and Top Listeners (hosts that receive the most traffic) for different traffic classes. If traffic in one of the Default or numbered port classes is regularly consuming a significant portion of total bandwidth (such as more than 10%), you can use the Top Talkers/Listeners feature to create new classes, thereby reducing the amount of unclassified traffic. You can turn on Top Talkers and/or Talk Listeners for a Default or numbered port class, and after a few days, look at a report to see if any single host is using a substantial percentage of the total (for example, more than 15%). You can then create a class for this host.

When PacketShaper is configured as a BCAAA client, it can retrieve user information from Microsoft Active Directory; user and group names can then appear in Top Talkers and Top Listeners reports.

To configure PacketShaper to track top hosts/users for a class:

1. Click the **Manage** tab.

2. In the left window pane, select the class for which you want to track hosts/users.

3. Select the **Top Talkers** and/or **Talk Listeners** checkboxes.

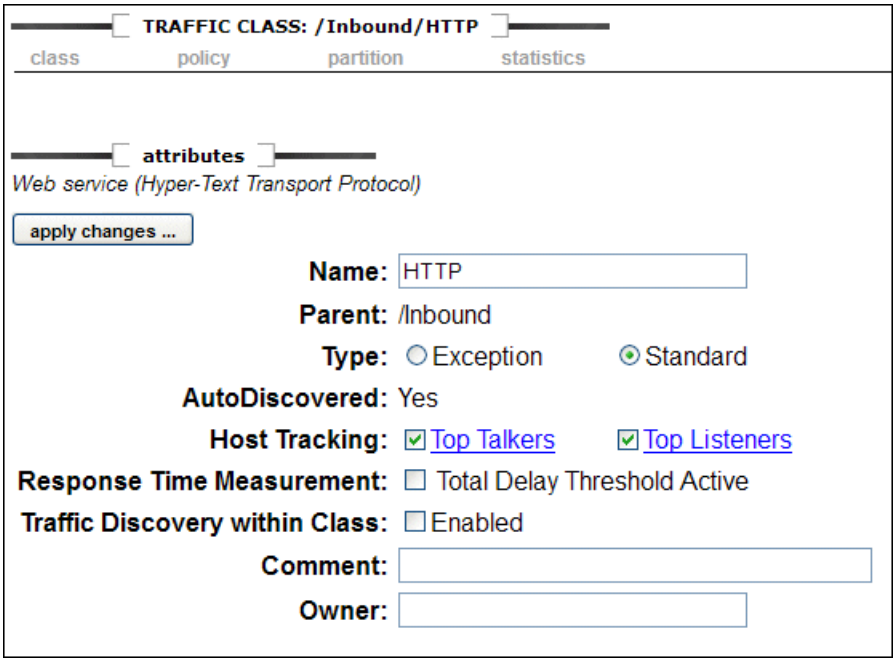

4. Click **apply changes**.

PacketShaper will now track the top twenty talkers/listeners for this class. See List Top Hosts for details on viewing a Top Talkers or Top Listeners report.

Note: You can enable a total of 32 Top Talkers and Top Listeners (combined) at one time. Thus, you can enable both Top Talkers and Top Listeners for up to 16 different classes. Or, if you enable one or the other (just Top Talkers or just Top Listeners), you can track top hosts on 32 different classes.

List Top Hosts or Users

After configuring PacketShaper to track the Top Talkers (hosts that initiate the most traffic) and Top Listeners (hosts that receive the most traffic) for a class, there are two ways to view the lists of top hosts: on the Monitor and the Manage tabs.

See Track Hosts and Users that Generate the Most Traffic if you haven't already enabled Top Talkers or Top Listeners for the class.

Displaying the List from the Monitor Tab

To access lists of Top Talkers and Top Listeners from the **Monitor** tab:

1. Click the **Monitor** tab.

2. In the *Monitor Traffic* window, select the appropriate icon in the **Top Host Analysis** column:

   - To list a class' Top Talkers, click the microphone icon.
   - To list a class' Top Listeners, click the ear icon.

   The *Top Talkers* screen displays the User Name, Group Name(s), DNS Name, IP Address, and Usage percentage for the class' 20 most active *sending* IPv4 and IPv6 hosts. The *Top Listeners* screen displays this information for the top *receiving* IPv4/IPv6 hosts.

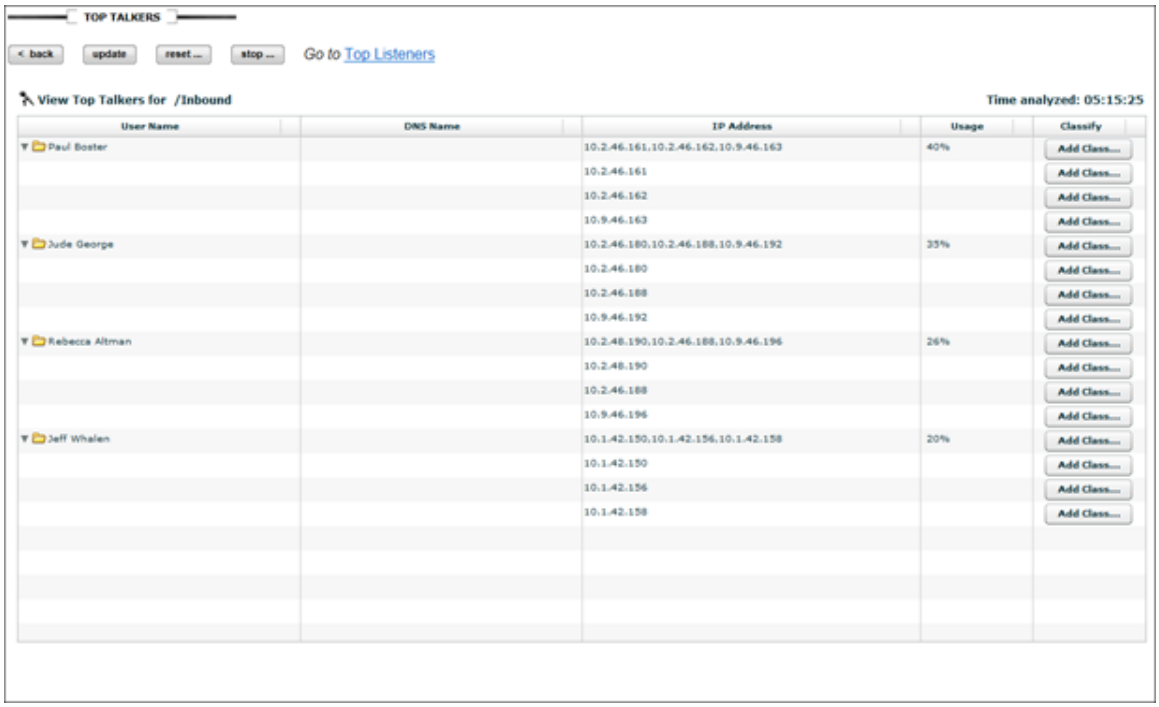

Additional Information

- If you see the message *User Awareness is not configured*, you have not configured PacketShaper to be a BCAAA client. In order to see traffic for Active Directory users, you must install and configure the BCAAA server.
- If you see the message *No DNS server address is configured*, you didn't specify a DNS server.

3. When a user belongs to multiple groups, the first group name is listed in the Group Name(s) column, followed by ellipses. To see all other groups associated with the user, mouse over the ellipses; the group names will pop up in a tool tip.

> **Note:** The only group names that display are ones for which a user group class has been created. The user may belong to other Active Directory groups, but they will not appear on the Top Talkers/Listeners report unless a class exists for that group.

4. When a user has multiple IP addresses assigned, the IP Address column lists all addresses, using elipses if they don't all completely fit within the column width. Click ▶ next to the user name to display each IP address on a separate line.

The Top Talkers/Listeners screen displays the hosts that have used the highest percentage of bandwidth in the class since tracking was enabled. A host stays on the top-20 list until another host uses more bandwidth, at which point the host may drop off the list entirely or move further down the list. For example, suppose top talkers is turned on for the Inbound/HTTP class, and cnn.com is the top consumer with 22%. If another host, yahoo.com, later consumes more bandwidth than cnn.com, yahoo.com might go to the top of the list and cnn.com would drop lower on the list.

Time analyzed indicates the length of time that PacketShaper has been tracking top hosts for the class. It is the amount of time since the Top Talkers/Top Listeners feature was enabled for the class, the **reset** button was clicked, or the unit was reset (whichever situation happened most recently).

The list is cleared when you click the **reset** button on this page or when you reset the unit.

- Use **update** to refresh the screen with the most current data.
- Use **reset** to clear the list and restart the host-tracking process.
- Use **stop** to disable tracking of Top Talkers or Top Listeners. (This is the equivalent to unchecking **Top Talkers** or **Top Listeners** on the **Manage** tab.)
- Click **Add Class** to quickly create a class for a specific user, DNS name, or IP address on the list. See Create a Class for a Top Host.

Displaying the List from the Manage Tab

Another way to list Top Talkers and Top Listeners is from the **Manage** tab:

1. Click the **Manage** tab.

2. In the left window pane, select the desired class.

3. In the *Traffic Class* window, click the **Top Talkers** or **Top Listeners** link. The list of top sending or receiving hosts is displayed.

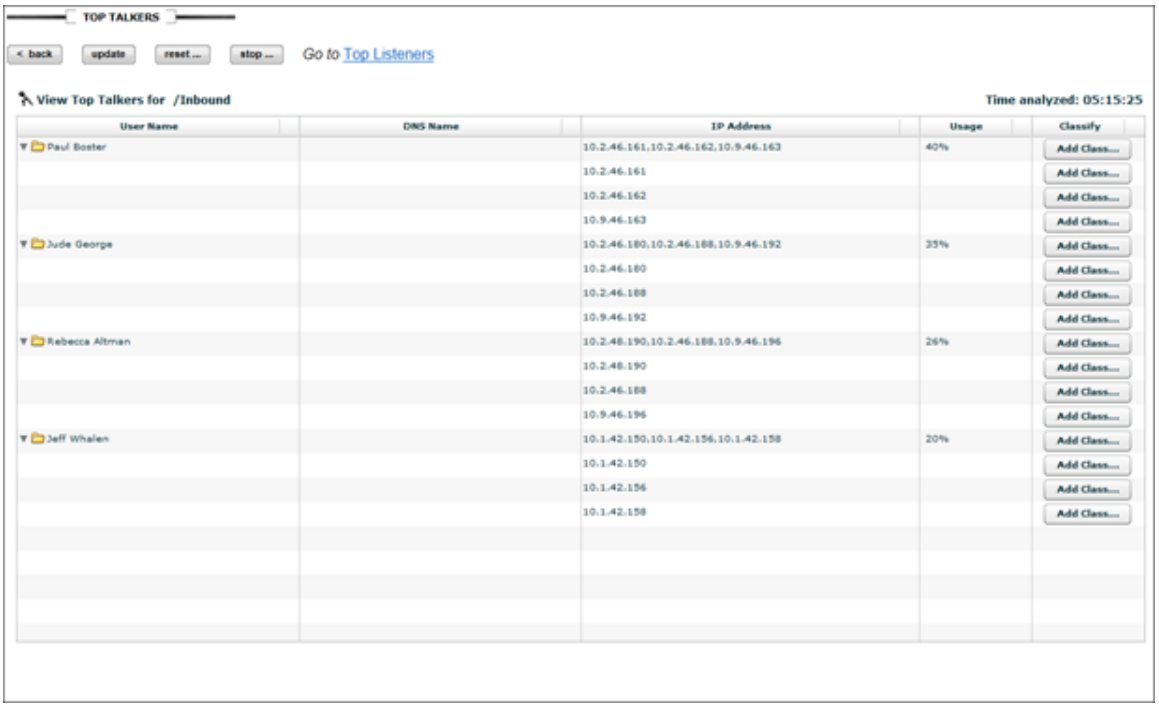

Create a Class for a Top Host or User

If you want to monitor or control the bandwidth of a top host or user, you can create a class for it. The *Top Talkers* and *Top Listeners* screens offer a quick way to do this.

To create a class for a top host or user:

1. Click the **Manage** tab.

2. In the left window pane, select the desired class.

3. In the *Traffic Class* window, click the **Top Talkers** or **Top Listeners** link. The list of top hosts and users displays.

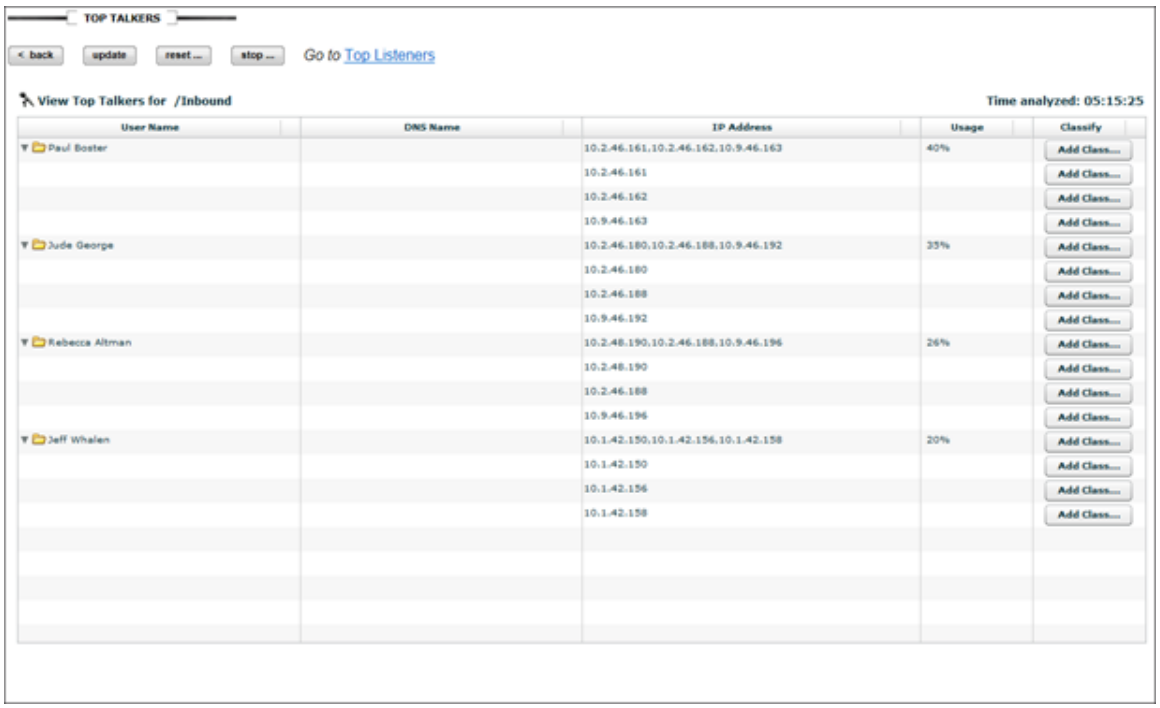

4. Click **Add Class** in the row of the host you want to classify. The *Top User Classify* window appears.

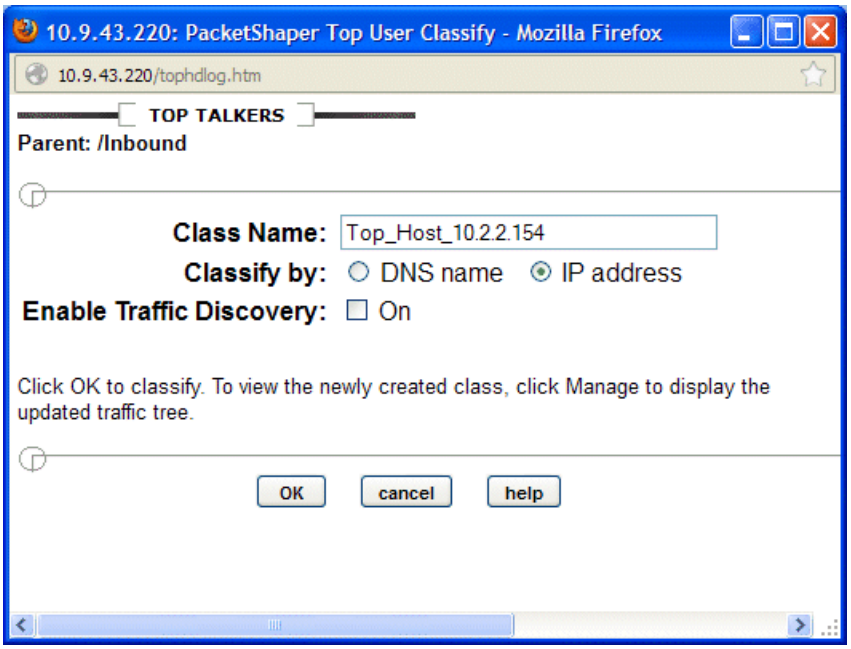

5. If the **Classify by** field is displayed, select the type of matching rule: **User name**, **DNS name**, or **IP address**.

> **Note:** If PacketShaper has identified the host by only one characteristic (user name, DNS name, or IP address), you will not be presented with **Classify by** options; it will automatically use the identified characteristic for the matching rule. However, if the DNS name and IP address (or a user name and IP address) are available for a host, you have the opportunity to select which matching rule you want to use.

6. In the **Class Name** field, modify the default name, if desired. For example, if the host is a server for a specific type of application, the class name could be the name of the application. The default name begins with *Top_Host* followed by the user name, DNS name, or IP address; for example *Top_Host_10.2.2.10*. If the user name contains UTF-8 characters, the **Class Name** field is blank.

> **Note:** Class names must be alphanumeric characters, up to 31 characters; hyphens, underscores, and periods are acceptable. Spaces and UTF-8 characters are *not* allowed in class names.

7. If you want PacketShaper to discover child classes of traffic in this new class, select **Enable Traffic Discovery**.

8. Click **OK**.

9. To view the newly created class, click the **Manage** or **Monitor** tab.

The new class will be a child of the class that is tracking top hosts.

Additional Information

- It's a good idea to check whether other IP addresses are associated with the application for which you are creating the class. If there are additional servers, you can add matching rules to define their IP addresses.
- When you classify by user name, PacketShaper will automatically classify all IP addresses associated with the user. For example, if a user has multiple devices, each with a different IP address, all the user's traffic will be classified into the user-based class.

Verify Traffic Classes

To verify that a traffic class is working as expected, go to the *Monitor Traffic* screen and observe the number of hits for that class, then start a new flow for the traffic class to see if the number of hits increases.

Different types of traffic interpret flows in different ways, so *hit counts* accrue at different rates. For example, a web browser may start one flow to get the HTML header for checking against its cache, a second flow to get the text/HTML portion of the page, and then several more flows to get each of the graphic elements on that page.

> **Note:** Instantaneous measurements tend to vary more than longer-term averages. It is common to see instantaneous spikes and dips in bandwidth, even though the one-minute average shows the flow maintaining a steady

Partition Tasks

One way to solve network and application performance problems is to create partitions. A partition manages bandwidth for a traffic class s aggregate flows, so that all of the flows for the class are controlled together as one. You can use partitions to protect mission-critical traffic by guaranteeing that a traffic class always gets a defined amount of bandwidth. In addition, you can use partitions to limit aggressive, non-critical traffic by allowing that traffic class to consume only a defined amount of bandwidth. For more information, see the Partition Overview.

This section includes tasks for creating, deleting, and modifying static and dynamic partitions.

Create a Partition

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a partition:

1. Click the **Manage** tab.

2. In the left window pane, select the class for which you want to create a partition.

3. In the right window pane, click **partition** and select **add**. The *New Partition* window displays. **show screen...**

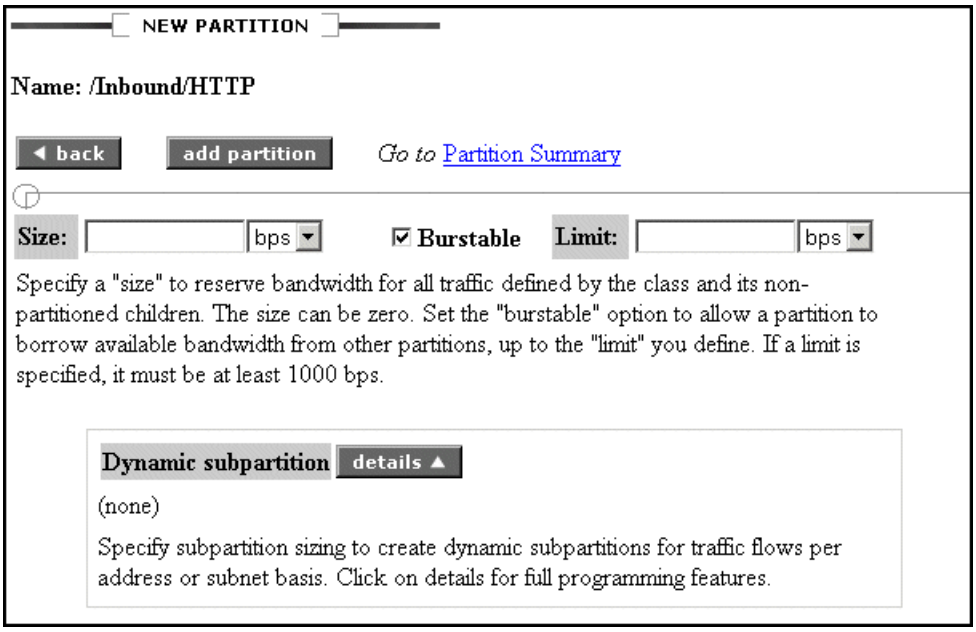

4. Define the partition using the following parameters:

Field	Description
Size	The minimum partition size, expressed in bits per second (**bps**) or as a percentage (**%**) of the parent partition's size. If the size is entered in bits per second, you can follow the integer by a k (thousands), M (millions), or G (billions). The minimum partition size is 1000 bps. The maximum partition size varies by model: PS-S200: licensed bandwidth PS-S400: 2 Gbps or licensed bandwidth (whichever is less) PS-S500: 2 Gbps. Exception: The Inbound and Outbound partitions can have a maximum size of 10 Gbps if you have a 10G license key or 5 Gbps if you have a 5G key. See Sizing a Static Partition for additional details and examples. **Note:** A guaranteed rate must be available within the minimum partition size. For the purpose of calculating whether there is enough bandwidth available in a partition to satisfy a guaranteed flow, it is assumed that all guaranteed rates are being used. For example, if the partition minimum size is 50k and the class has a policy with a guaranteed rate of 20k, a maximum of two flows can exist, even if the two flows are only using 15k.
Burstable	When selected, allows the partition to use available excess bandwidth if needed. When this box is unchecked, the partition has a fixed size; when it is not using its reserved bandwidth, that bandwidth is available to other traffic. **Note:** A burstable partition simply allows *larger* flows, not *more* flows.
Limit	A maximum bandwidth limit to be used when the partition bursts. This limit must be greater than the partition minimum. The limit can be specified in bits per second or as a percentage of the parent partition's size.

5. Click **add partition** to create the partition. If you specified the partition size as a percentage, PacketShaper will calculate the absolute size in bits per second and display the results in the *Traffic Class* window.

6. If you want to make this partition dynamic, click **details** and then refer to Create a Dynamic Partition.

Note: In order for partitions to take effect, traffic shaping must be enabled. See Enable/Disable Traffic Shaping.

Tip: When creating partitions, make sure you don't allocate bandwidth in such a way that Inbound/Default and Outbound/Default get starved –that is, there is no bandwidth available for these classes. If this happens, traffic classification and policies may not work as expected.

Sizing a Static Partition

Determining an appropriate minimum and maximum size for a static partition is a learned skill. These guidelines can help.

The sections contained below are:

What is a Static Partition?
Choosing a Unit of Measurement (bps or %)
Partition Examples
Is the Partition Size Appropriate?

What is a Static Partition?

A static partition defines the amount bandwidth that all traffic flows in one traffic class (including all child or subordinate classes) can access. A static partition does *not* control how total bandwidth is portioned out to each matching flow.

Each *burstable* static partition has two associated sizes (Size and Limit); each *non-burstable* (or *fixed*) partition has one associated size (Size):

1. **Size**: the minimum amount of bandwidth the traffic class can dependably access, no matter how much bandwidth contention exists in other classes. If the traffic class does not need some or all of its minimum size, then PacketShaper loans the bandwidth to other traffic in need.

2. **Limit**: the maximum amount of bandwidth the traffic class can ever access, no matter how much unused, excess bandwidth is available. If a traffic class is already using its minimum size of bandwidth and still could use more, PacketShaper calculates additional allocation based on the amount of available bandwidth and the priorities of competing traffic.

Choosing a Unit of Measurement

When you define the minimum and maximum sizes, you also choose your unit of measurement: BPS or %. You can specify either an *explicit* size with a bits-per-second value (100K, for example), or you can specify a *relative* size with a percentage of a parent partition (30%, for example). Each of these two methods have appropriate and inappropriate contexts for use.

Not all classes have partitions, so when you specify a partition's size as a percentage, note that it will be a percent of the nearest parent class with a partition. The nearest parent class with a partition could be /Inbound or /Outbound, or it could be an intermediate class if your tree has multiple levels of hierarchy.

When to Use Explicit Partition Sizes

Types of situations most suited for explicit, bits-per-second partition sizes:

- Typical "bandwidth farming" situations your focus is on allocating prescribed amounts of bandwidth. Examples:
  - A service provider with per-subscriber contractual obligations
  - A university with planned bandwidth amounts for each dormitory
  - A building manager giving set amounts of bandwidth to each office or apartment

- You want to accommodate a set number of users for an application with predictable per-flow behavior. If you know that each user will use up to X Kbps, and you want to accommodate Y users, then you can define a partition with a size of X*Y Kbps. These cases usually involve streaming applications, and a rate policy can keep each stream at the bandwidth rate it needs for smooth performance and predictability. Examples:
  - Voice over IP
  - Internet radio
  - Distance learning

- You want to avoid an absolute prohibition against an application, but you don't want it to get very much bandwidth. Examples:
  - Sites with questionable content
  - Recreational applications or games in a business setting

When to Use Relative Partition Sizes

Types of situations most suited for relative, percentage partition sizes:

- You have a primary WAN link and a back-up WAN link with different capacities and you don't want to change partition configurations when your backup becomes active.

- You have a mobile network with a WAN link provided by satellite; your WAN capacity varies depending on which satellite you're near.

- Your focus is on application performance, rather than fulfilling contractual obligations for amounts of bandwidth.

- Your performance strategy involves dividing up WAN capacity by application, and you are not dealing with streaming applications and set numbers of users. You simply want to divide your bandwidth pie proportionately, according to business and application requirements.

Partition Examples

Partition for Music Downloads

The size of your partition should vary depending on how restrictive you want to be and the relative importance of other traffic. Typically you want to allow music downloads, but keep them from disrupting other traffic. A good starting point might be to reserve anywhere from 0 to 5 percent of your link, select burstable, and cap the combined size of inbound and outbound partitions at 5 to 25 percent of your link capacity.

If you want to use explicit sizes, then suppose you want to reserve 5 percent of your T1 link for P2P traffic and cap it at 20 percent. A T1 is about 1.5 Mbps (bi-directional or full duplex). Five percent of each direction is about 75 Kbps, and 20 percent is about 300 Kbps.

Partition for Citrix

The size should be at least sufficient to cover the number of users you expect with minimally acceptable performance. For example, if you expect to have 20 concurrent users who each need at least 15 Kbps for minimally acceptable performance, then make sure your partition has a minimum size of 300 Kbps.

Partition(s) for Streaming Media

The minimum and maximum sizes of your partitions depend on how restrictive you want to be and the relative importance of streaming media with other traffic.

For example, suppose you want to ensure good performance for a streaming application that requires a dedicated 15 Kbps for each active session. You want to make sure there's room for at least five sessions at all times. And you want to accommodate up to 20 sessions. You have a T1 link, which is about 1.5 Mbps and bi-directional or full duplex. Multiplying 5 sessions by 15 Kbps gives a minimum of 75 Kbps, or about 5 percent of your capacity. Multiplying 20 sessions by 15 Kbps gives a maximum of 300 Kbps, or about 20 percent of capacity. Your partition definition would be: Size: 75K, burstable, Limit: 300K.

Alternatively, you might not want to worry about per-session performance, as streaming media is not critical for your organization. You only know that you're willing to devote a maximum of 10 percent of your link capacity to streaming media in order to protect bandwidth for other applications. With your same T1 link, you could define a burstable partition with a minimum size of 0 Kbps and a maximum of 10 percent (assuming that the nearest parent with a partition is /Inbound or /Outbound).

Is the Partition Size Appropriate?

To analyze whether a partition size is appropriate, you can use the following method:

Create Partition Utilization Graphs

Modify a Partition

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To modify the settings of an existing partition:

1. Click the **Manage** tab. The *Traffic Class* window appears.

2. In the left window pane, click the class name.

3. Click **partition** in the *Traffic Class* window and select **open**, or click the icon to the right of the class name in the traffic class tree. The *Partition* window appears. **show screen...**

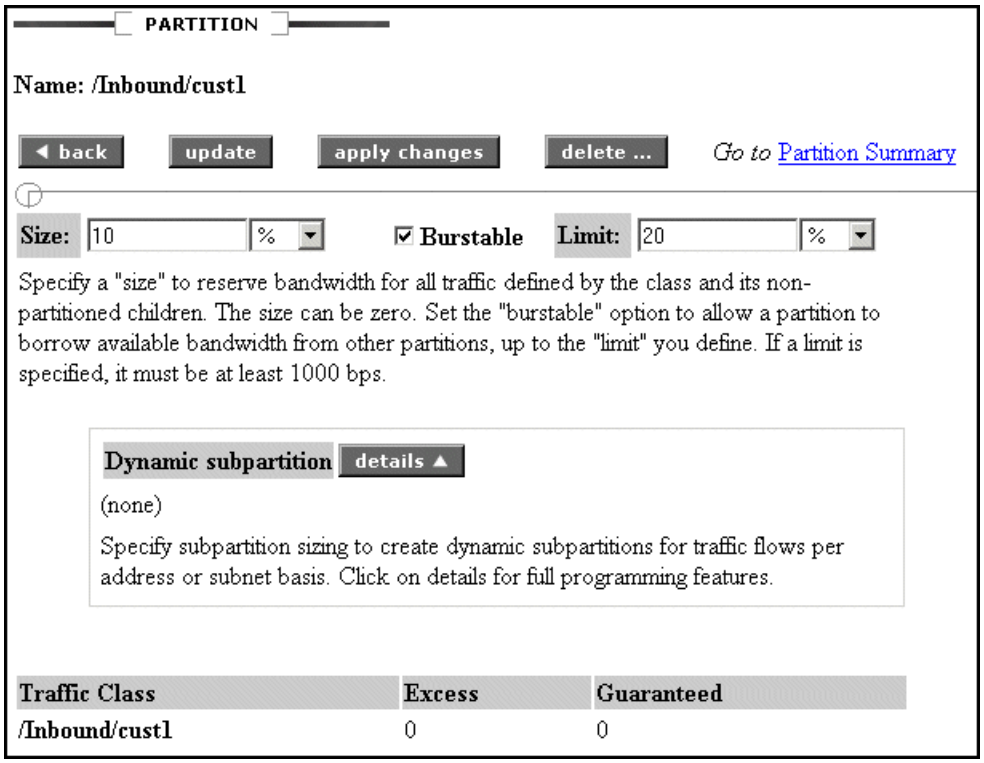

4. Update the partition's size, burstable, and limit parameters (described in Create a Partition).

5. Click **apply changes** to update the partition.

6. Click **refresh** to display the updated list of partitions in the *Partition* window.

> **Tip:** Another way to open the *Partition* window is to go to the **Monitor** tab and, in the Partition Min-Max column, click the setting you want to modify.

View Partition Information

You can view information about your partitions in two places: the *Monitor Traffic* window and the *Partition Summary* window.

List Classes with Partitions

The *Monitor Traffic* window can be set to show only the classes that have partitions. To display this list:

1. Click the **Monitor** tab. The *Monitor Traffic* window appears. **show screen...**

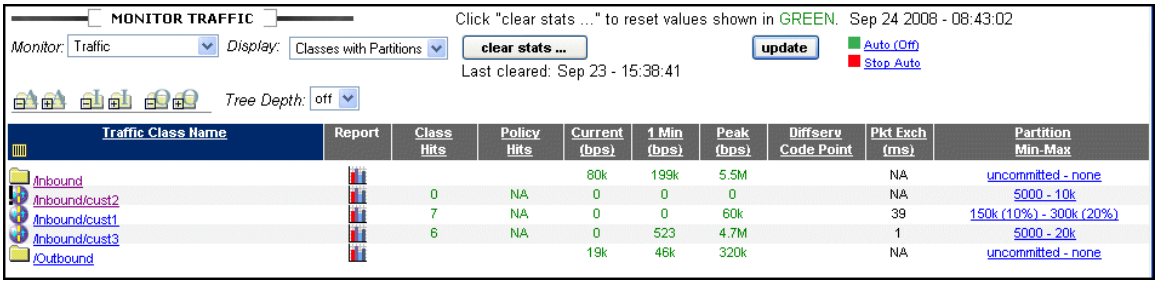

2. In the **Display** drop-down list, select **Classes with Partitions**. The window now lists only the classes that have partitions. **show screen...**

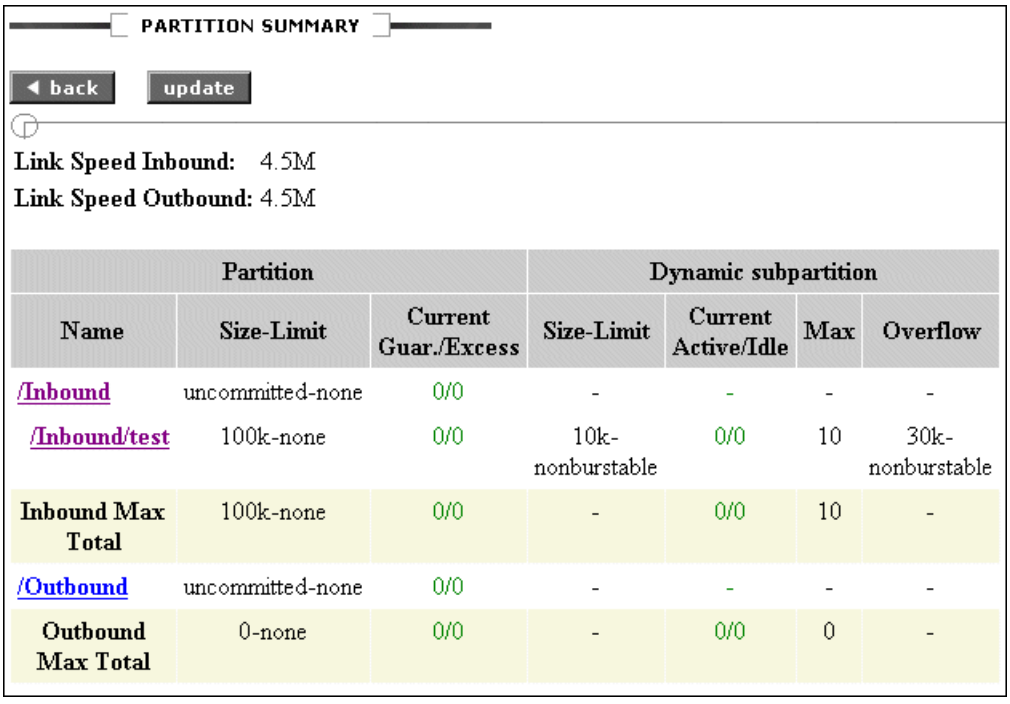

The Partition Min-Max column defines the amount of bandwidth allocated for the class:

- *Uncommitted* indicates that the guaranteed minimum allocation is whatever is not committed to other partitions. Normally uncommitted is used only by the default Inbound and Outbound partitions.
- *None* indicates that the partition is burstable without a limit.
- If the minimum and maximum values are the same, the partition is fixed and non-burstable.
- If a down arrow appears next to the partition size, the partition is dynamic. Place the mouse pointer on the down arrow to display the parameters of the dynamic partition: minimum and maximum subpartition size, the number of active and idle users currently using subpartitions, and the maximum number of subpartitions available.
- An asterisk following a partition's minimum size indicates that the configured minimum value has been scaled to the displayed value, due to the link being overcommitted.
- If a percentage appears after the partition size, the partition was specified as a percentage of the parent partition's size. For example, *450k (10%)* for the Inbound/HTTP class indicates that HTTP's partition size is 10% of the Inbound partition's size, resulting in a partition size of 450 Kbps (assuming the link size is 4.5 Mbps).

View the Partition Summary

The *Partition Summary* window provides an overview of the current partitions. To view the *Partition Summary* window:

1. Click the **Manage** tab. The *Traffic Class* window appears.
2. Click **partition summary** above the traffic tree. The *Partition Summary* window appears. **show screen...**

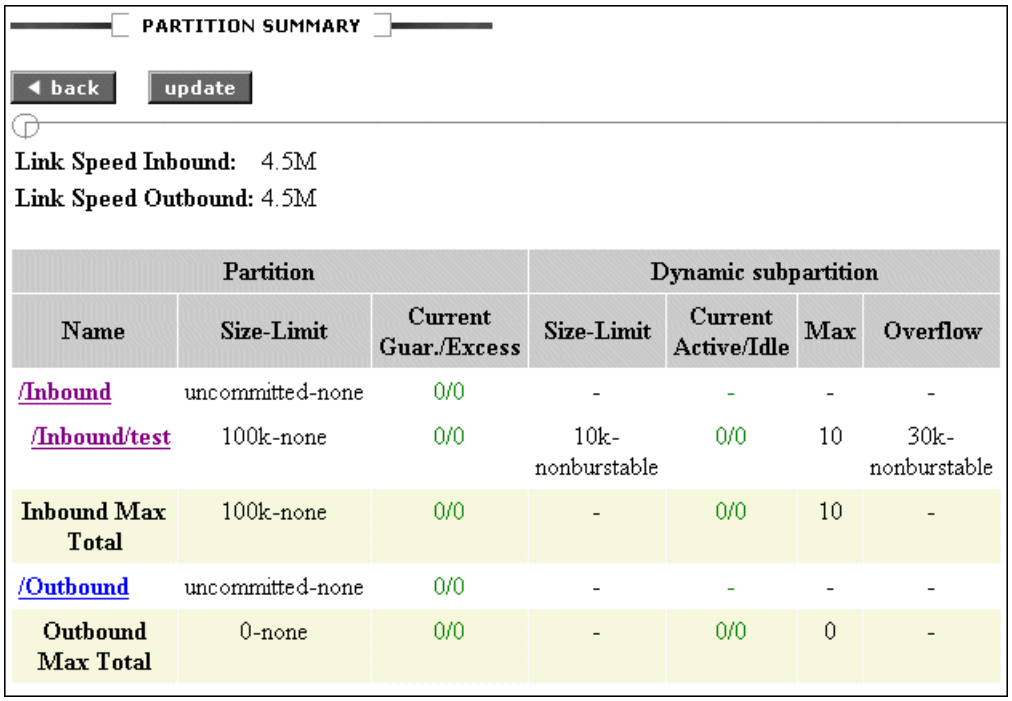

Partition			Dynamic subpartition			
Name	Size-Limit	Current Guar./Excess	Size-Limit	Current Active/Idle	Max	Overflow
/Inbound	uncommitted-none	0/0	-	-	-	-
/Inbound/test	100k-none	0/0	10k-nonburstable	0/0	10	30k-nonburstable
Inbound Max Total	100k-none	0/0	-	0/0	10	-
/Outbound	uncommitted-none	0/0	-	-	-	-
Outbound Max Total	0-none	0/0	-	0/0	0	-

> **Tip:** Another way to display the *Partition Summary* window is to click the **Partition Summary** link in the *Partition* window.

The *Partition Summary* window displays the following information:

Field	Description
Link Speed	The access link rate for the /Inbound and /Outbound partitions
Name	The name of the partitioned traffic class
Partition Size	The partition size in bits per second. If a percentage appears after the partition size, the partition was specified as a percentage of the parent partition's size. Example: *450k (10%)* for the Inbound/HTTP class indicates that HTTP's partition size is 10% of the Inbound partition's size, resulting in a partition size of 450 Kbps (assuming the link size is 4.5 Mbps). An *uncommitted* rate is the difference between the current link speed and the sum of all committed partition minimums. *Uncommitted* indicates that: The size is dynamically determined at run time. The minimum allocation is whatever is not committed to other partitions. If a partition is oversubscribed, an asterisk appears after the partition's minimum size value, indicating that the configured minimum value has been scaled to the displayed value. The percentage overcommitted is displayed in red for example, Overcommitted by 1%.
Partition Limit	The maximum bandwidth allowed for the partition (when excess bandwidth is available)
Partition Current Guar.	The guaranteed rate in bits per second that is in use > **Note:** This column is only relevant if the traffic managed under the partition is also managed by a rate policy. The portion of the bandwidth guaranteed in the rate policy that is currently being used is listed. For more information, see Set a Rate Policy.
Partition Current Excess	The amount of excess rate in bits per second currently in use (not including the current guaranteed bandwidth)
Dynamic subpartition Size	The minimum amount of bandwidth assigned to each user, expressed in bits per second. If a percentage appears after the size, the subpartition was specified as a percentage of the static partition's size.
Dynamic subpartition Limit	When excess bandwidth is available, the maximum amount of bandwidth each user can access

Field	Description
Dynamic subpartition Current Active	The number of subpartitions in active use
Dynamic subpartition Current Idle	The number of subpartitions that were created but have not been active for 300 seconds (5 minutes). Idle subpartitions still have flows that are sending packets.
Dynamic subpartition Max	The maximum number of subpartitions
Dynamic subpartition Overflow	The minimum and maximum amount of bandwidth allotted to an overflow subpartition (used when the maximum number of subpartitions has been reached)

Create Partition Utilization Graphs

One way to check whether a partition is the proper size is to create partition utilization graphs, to compare usage over a period of time with the defined partition size. The *Partition Utilization* graph shows a partition's average bandwidth usage in bits per second, with horizontal lines indicating the partition's minimum and burstable size limits. The *Partition Utilization with Peaks* graph shows a partition's average and peak bandwidth usage. With both of these graphs, it's recommended that you select the checkbox **Include Link and Partition Sizes in Rate Graphs** so that you can easily compare actual usage with the partition size. If traffic is frequently peaking at the partition maximum, you may want to consider increasing the size of the partition (assuming it is traffic you are trying to *protect*, not *contain*).

To create a partition utilization report:

1. In the **Manage** tab:

   - In the left window pane, select the partition you want to graph.
   - Click **statistics** and choose **graph**. The *Statistics:Reports* screen appears.

2. In the Graph Type column, select two graphs: Partition Utilization and Partition Utilization with Peaks.

3. For each graph, enter a number in the **Period** column and select the units (**hour**, **day**, **week**, or **month**). The default is **1 hour**.

4. Select an end date for each graph. The default end date is **now** (the current date).

To change the date, select the month and day from the drop-down lists. You can change the date to any time within the past two months.

5. Select an end time for each graph. The default time is **now** (the current time).

To change the time, select the time from the drop-down list.

6. Select the **Include Link and Partition Sizes in Rate Graphs** checkbox. When this box is checked, the partition size will appear as a horizontal line at the top of partition utilization graphs.

7. Click **create report** to generate the report.

In the graph below, you can see that the traffic is periodically peaking near the partition's burst limit. If this was mission-critical traffic you were trying to protect, you might consider enlarging the partition size.

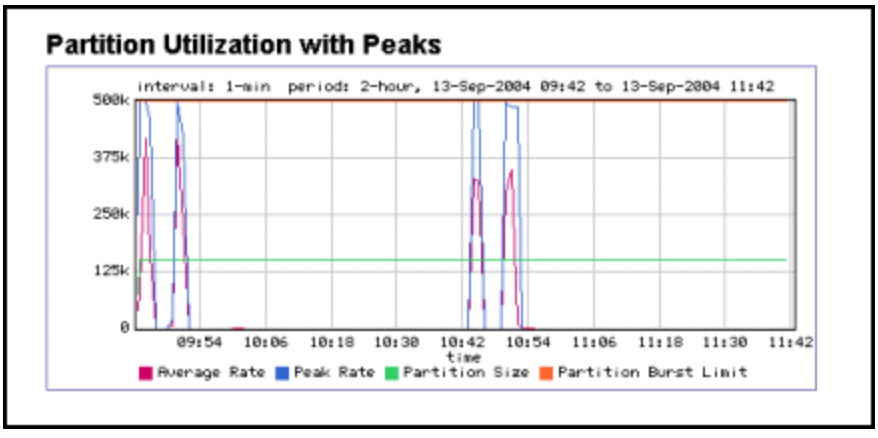

Delete a Partition

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To delete a partition:

1. Click the **Manage** tab. The *Traffic Class* window appears. **show screen...**

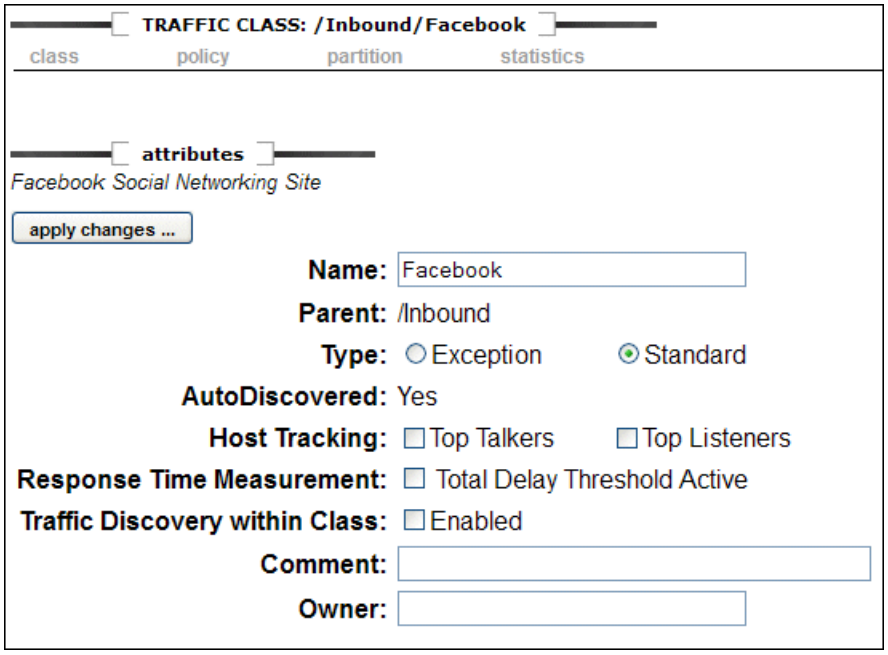

2. In the left window pane, click the class name.
3. Click **partition** in the *Traffic Class* window, or click the icon to the right of the class name in the traffic class tree.
4. Click **delete**.
5. Click **OK** in the confirmation box to remove the partition.

Create a Dynamic Partition

With PacketShaper's dynamic partition feature, subpartitions are automatically created as users become active in a traffic class. This capability allows service providers or enterprise customers to provision precise amounts of bandwidth on a per-user basis in just one step. Dynamic partitions also can be used to guarantee fairness per subpartition even in configurations where bandwidth is oversubscribed. So, if demand ever exceeds available bandwidth, dynamic partitions can ensure that each user will get the same amount of bandwidth.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To make a partition dynamic:

> **Note:** To create a dynamic partition, you must first create a static partition for the traffic class and then specify the dynamic partition settings. If you already have a static partition defined for the traffic class, open the partition definition (see Modify a Partition) and then specify the dynamic partition settings.

1. In the *New Partition* (or *Partition*) window, click **details**. The *Dynamic Subpartition* window appears. **shown screen...**

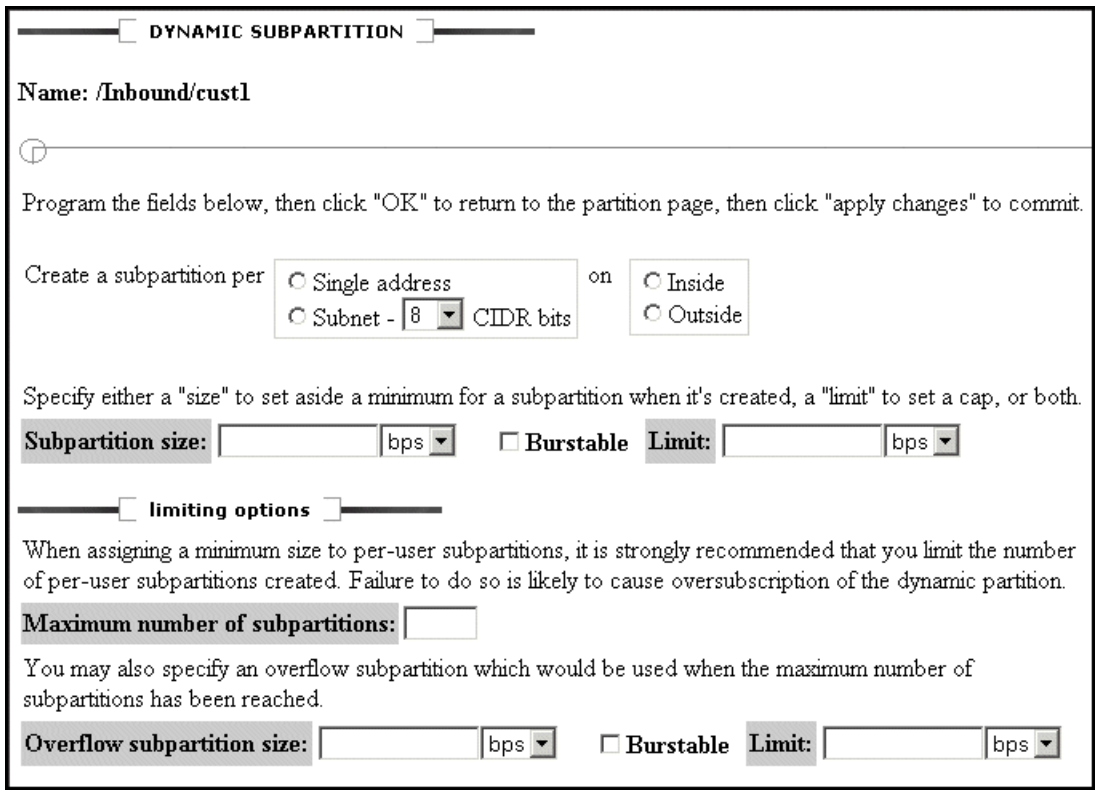

2. Define the user as a single IP address or a subnet:

- To define a user as a single IP address, select **Single address**.
- To define a user as a subnet, select **Subnet**, then select the CIDR value that represents the number of constant bits in the address range.

3. Select the side (**Inside** or **Outside**) on which the different IP addresses or subnets that will generate dynamic subpartitions will be found.

For example, in a configuration in which the PacketShaper unit sits between a LAN and the Internet, select **Inside** to create a dynamic subpartition for each inside client (LAN user) that connects to any outside host. Alternatively, select **Outside** to create a subpartition for each outside host connected to any inside client.

> **Note:** Dynamic subpartitions can be based on inside or outside systems, independent of the class' matching rules (which may specify addresses or subnets specifically on the inside or outside). This allows subpartitioning based on the addresses you classify, or their communicating peer systems.

4. Define the new subpartition using the following parameters:

Field	Description
	The minimum amount of bandwidth to be assigned to each user, expressed in bits per second (**bps**) or as a percentage of the static partition's size (**%**). If the size is entered in bits per second, you can follow the integer by a k (thousands) or M (millions). The minimum subpartition size is 1000 bps.
Subpartition size	Set this field to zero to have PacketShaper allocate bandwidth equitably to each subpartition, so that the total of all subpartitions equals the static partition's size. **Note:** Minimum subpartition size is frequently best handled by setting this field to zero and setting a maximum number of subpartitions (covered below). However, you must use a non-zero size if you want to implement per-session guaranteed rates within rate policies for this same traffic.
Burstable	When selected, allows each subpartition to use available excess bandwidth if needed; each subpartition will be allocated the same amount of excess bandwidth. If you set a subpartition size value of zero, you must enable the **Burstable** check box. When the subpartition's size is explicit (not 0) and this box is unchecked, each subpartition has a fixed size; when it is not using its reserved bandwidth, that bandwidth is available to other traffic.
Limit	The maximum amount of bandwidth to be assigned to each subpartition. The limit can be specified in bits per second or as a percentage of the static partition's size. Use this field only if you want to enforce a cap on each user or subnet even if more bandwidth is available. Managed bandwidth service providers are most frequently in this position, needing to cut off usage at agreed-upon, paid-for limits. If you don't want to enforce a cap, leave the field blank. Even if this field is left blank, the limit on the static, parent partition still restricts the total bandwidth for the aggregate of all subpartitions.

5. *Optional:* Set a cap on the number of active users in the dynamic partition by entering a value in the **maximum number of subpartitions** box.

> **Note:** If the partition is fixed (non-burstable), be sure to put a cap on the dynamic subpartitions and create an overflow partition. This will prevent oversubscription problems.

6. *Optional:* Create an overflow partition for users to tap into if the cap is exceeded. Select the **Enable** checkbox and then use the following parameters:

Field	Description
Overflow subpartition size	The minimum amount of bandwidth in the overflow partition, expressed in bits per second (**bps**) or as a percentage of the static partition's size (**%**). The minimum overflow subpartition size is 1024 bps. If you leave this field blank, the overflow subpartition's minimum size will be *uncommitted* (whatever is not committed to other partitions).
Burstable	When selected, allows the overflow partition to use available excess bandwidth if needed. When this box is unchecked, the overflow partition has a fixed size; when it is not using its reserved bandwidth, that bandwidth is available to other traffic.
Limit	The maximum amount of bandwidth in the overflow subpartition. The limit can be specified in bits per second or as a percentage. Must be blank or a number larger than the size field.

> **Note:** If you don't create an overflow partition, flows will be refused when the cap is reached or when the partition limit on the box is exceeded.

7. Click **OK**.

8. Click **add partition** to create the dynamic subpartition.

Examples of dynamic partitions and their effects:

Dynamic Partition Values	Effects
Static partition Size: 1.5 Mbps Burstable: not selected Limit: (blank)	With two users: they share the entire 1.5 Mbps. If one uses less than 750 Kbps, then the other can access more. If the two users don't use the 1.5 Mbps, it's available to other traffic classes outside the static partition.
Dynamic subpartitions Subpartition size: 0 Burstable: checked Limit: (blank) Maximum #: (blank) Overflow partition: (blank)	With 50 users: each gets 30 Kbps. If any don't use their quota, it's shared among others that can use it. If there's still leftover bandwidth in the static partition after all 50 users are serviced, it's available to other traffic classes. With 100 users: each gets 15 Kbps. Nobody can use more unless someone uses less. With 200 users: each gets 7.5 Kbps.

Dynamic Partition Values	Effects
Static partition Size: 1.5 Mbps Burstable: not selected Limit: (blank) **Dynamic subpartitions** Subpartition size: 0 Burstable: checked Limit: (blank) Maximum #: 100 Overflow partition: (blank)	With two users: they share the entire 1.5 Mbps. If one uses less than 750 Kbps, then the other can access more. If the two users don't use the 1.5 Mbps, it's available to other traffic classes outside the static partition. With 50 users: each gets 30 Kbps. If any don't use their quota, it's shared among others that can use it. If there's still leftover bandwidth in the static partition after all 50 users are serviced, it's available to other traffic classes. With 100 users: each gets 15 Kbps. Nobody can use more unless someone uses less. With 200 users: 100 each get 15 Kbps, 100 get refused.
Static partition Size: 1.5 Mbps Burstable: not selected Limit: (blank) **Dynamic subpartitions** Subpartition size: 0 Burstable: checked Limit: 30 Kbps Maximum #: 100 Overflow partition: 100 Kbps	With two users: each have 30 Kbps. The remaining 1.4+ Mbps is available to other traffic classes outside the static partition. With 50 users: each gets 30 Kbps. If any don't use their quota, it's *not* shared among other subpartition users. If there's leftover bandwidth, it's available to other traffic classes outside the static partition. With 100 users: each gets 15 Kbps. Nobody can use more unless someone uses less. And nobody can use more than 30 Kbps. With 200 users: 100 each get 14 Kbps. And the remaining 100 share the overflow partition at 1 Kbps each.

After the dynamic partition is set up, whenever a new user begins generating flows in that class, a subpartition will be created for the user on the fly. The per-user partition remains in existence until it's reused for new flows by the same user or needed by another user.

Modify a Dynamic Partition

When you modify dynamic partition settings, keep the following in mind:

- If you modify the partition's structural configuration (such as the cap or per-subnet/per-address setting), any existing subpartitions are *not* affected by the change. When new users become active, the new subpartitions will utilize the modified settings (and eventually the old subpartitions will be replaced by new ones with the new settings).
- If you modify the size of the subpartition or the overflow partition, the new sizes are immediately applied to existing subpartitions.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To modify the settings of an existing dynamic partition:

1. Click the **Manage** tab. The *Traffic Class* window appears.

2. In the left window pane, click the class name.

3. Click **partition** in the *Traffic Class* window and select **open**, or click the icon to the right of the class name in the traffic class tree. The *Partition* window appears. **show screen...**

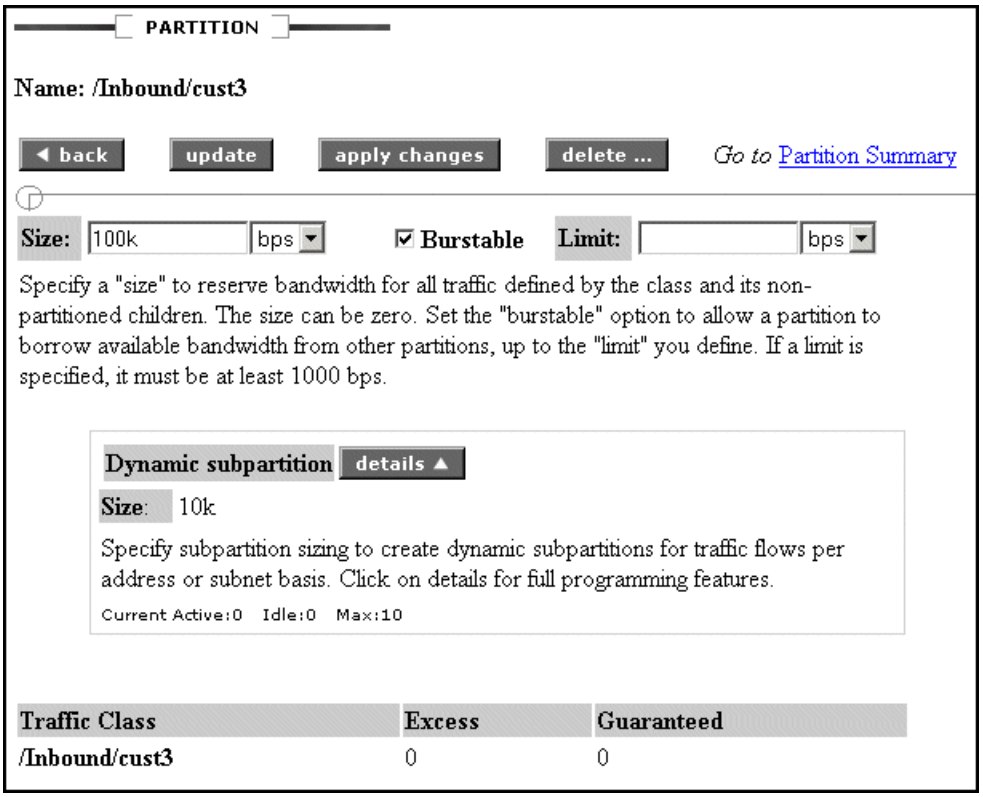

4. Update the partition's size, burstable, and limit parameters (described in Create a Partition).

5. To modify the subpartition parameters, click **details**. The *Dynamic Subpartition* window opens. **show screen...**

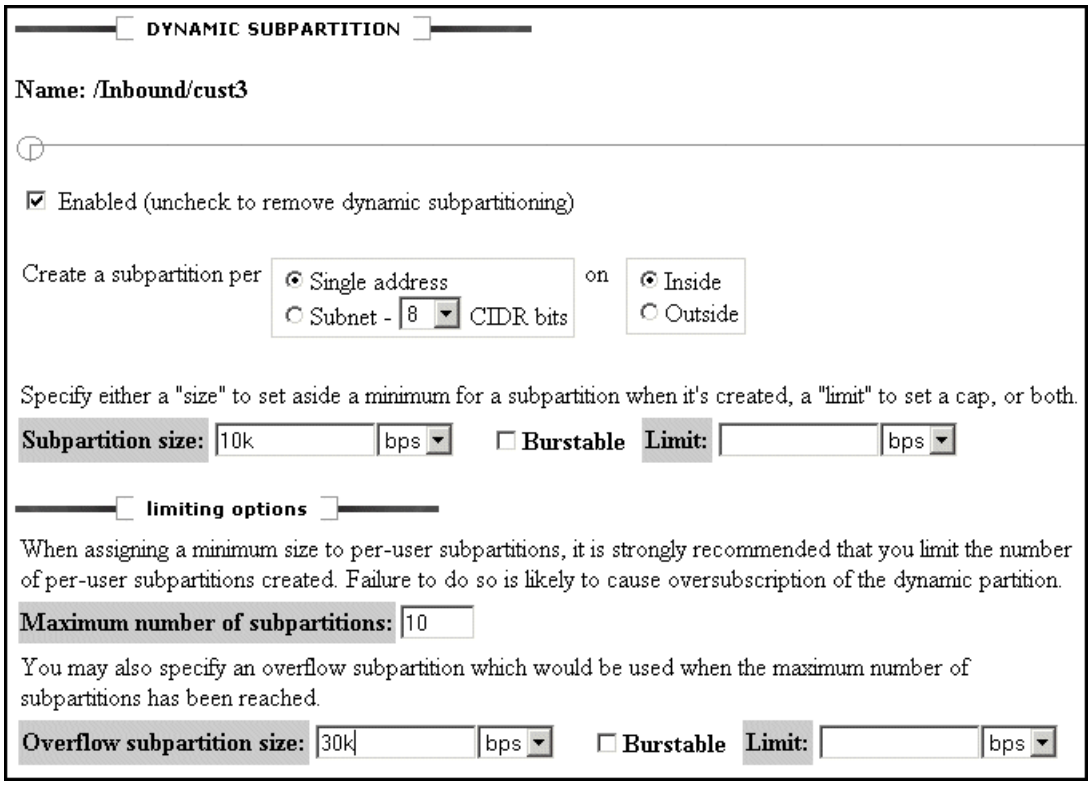

6. Update the subpartition settings (described in Create a Dynamic Partition) and click **OK**.

7. Click **apply changes** to update the partition.

8. Click **refresh** to display the updated list of partitions in the *Partition* window.

Remove a Dynamic Partition

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To remove a dynamic partition (but leave the static partition in place):

1. Click the **Manage** tab on the navigation bar. The *Traffic Class* window appears.

2. In the left window pane, click the class name.

3. Click **partition** in the *Traffic Class* window and select **open**, or click the icon to the right of the class name in the traffic class tree. The *Partition* window appears.

4. Click **details**. The *Dynamic Subpartition* window appears.

5. Uncheck the **Enabled** checkbox to remove dynamic subpartitioning of this class. **show screen...**

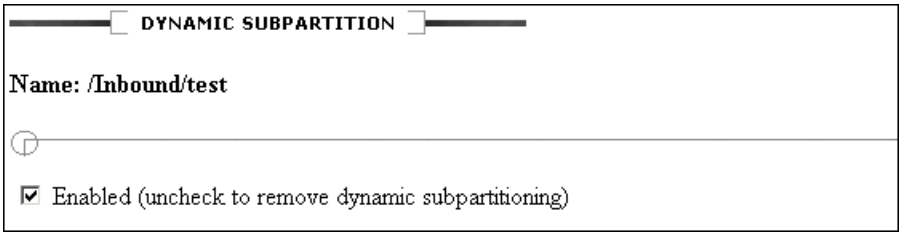

6. Click **OK**.

7. Click **apply changes** in the *Partition* window.

Enable/Disable Traffic Shaping

When shaping is on, traffic is classified and measured, and policies and partitions actively control bandwidth allocation. When shaping is off, traffic is classified and measured, but it is not actively managed.

> **Note:** Partitions and policies (with the exception of ignore policies) do not take effect until traffic shaping is enabled.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To turn traffic shaping on or off:

1. Click the **Setup** tab. The *Basic Settings* screen is automatically displayed in the *Setup* window. **show screen...**

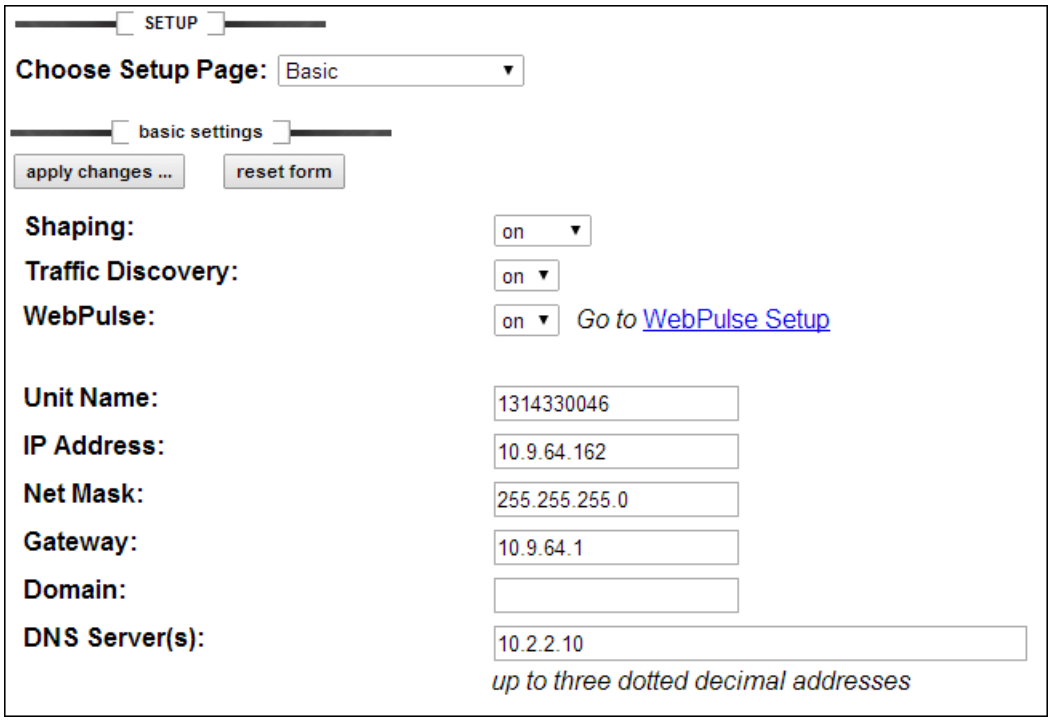

2. In the **Shaping** field, select **on** to enable shaping or **off** to disable.

3. Click **apply changes** to update the settings.

Policy Tasks

One way to improve network and application performance is to create policies. A policy determines how an application s individual flows are treated in the context of competing applications, and allows you to manage bandwidth on a flow-by-flow basis. With policies, you can give each flow of mission-critical traffic the bandwidth it needs for optimum performance, as well as protect it from greedy, less important traffic. In addition, policies can keep non-urgent traffic flows (such as FTP) from consuming more than an appropriate share of bandwidth. For more information, see the Policy Overview.

Use Suggested Policies

PacketShaper includes a set of suggested policies for many supported protocols and services. These policies were developed for a typical enterprise branch office and may or may not be suitable to your situation. You can apply a suggested policy to a traffic class, whether it was manually created or automatically created by the traffic discovery process.

> **Note:** All suggested rate policies set the guaranteed rate to zero with bursting enabled. If you have a specific minimum rate requirement for an application, such as Voice over IP, you will need to modify the suggested policy.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To create a policy using PacketShaper's suggested settings:

1. Click the **Manage** tab. The *Traffic Class* window appears.
2. Select a class in the tree in the left window pane.
3. Click **policy** in the *Traffic Class* window and select **add**.
4. Click **suggest policy** in the *New Policy* window. **show screen...**

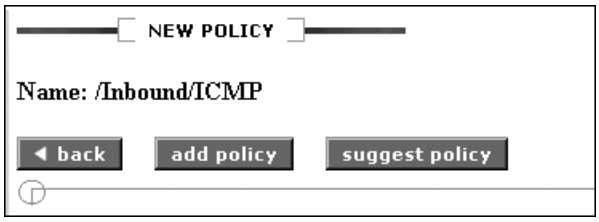

A pop-up window appears, notifying you whether a default policy was found for this service type. **show screen...**

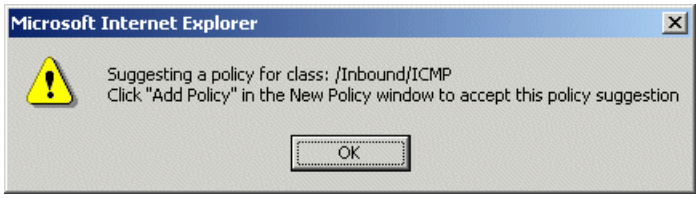

5. Click **OK** in the pop-up window.
6. To fine-tune the policy, overwrite the displayed values or select a new policy type.
7. To create the policy, click **add policy** in the *New Policy* window.

Set a Rate Policy in the Advanced UI

Using a *rate policy*, you can establish a minimum rate for each flow of an application, allow the flow prioritized access to excess bandwidth if it's available, and set a limit on the total bandwidth that a flow can consume. A rate policy is especially suited for traffic that tends to burst.

A rate policy enables precise management of TCP traffic. For User Datagram Protocol (UDP) traffic, a rate policy uses delay-bound scheduling. Rate policies are *not* recommended for non-IP traffic.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To apply a rate policy to a traffic class:

1. Click the **Manage** tab.

2. In the left window pane, select the class to which you want to apply a rate policy.

3. Click **policy** and select **add**. The *New Policy* window appears.

4. Click **rate** in the *New Policy* window. **show screen...**

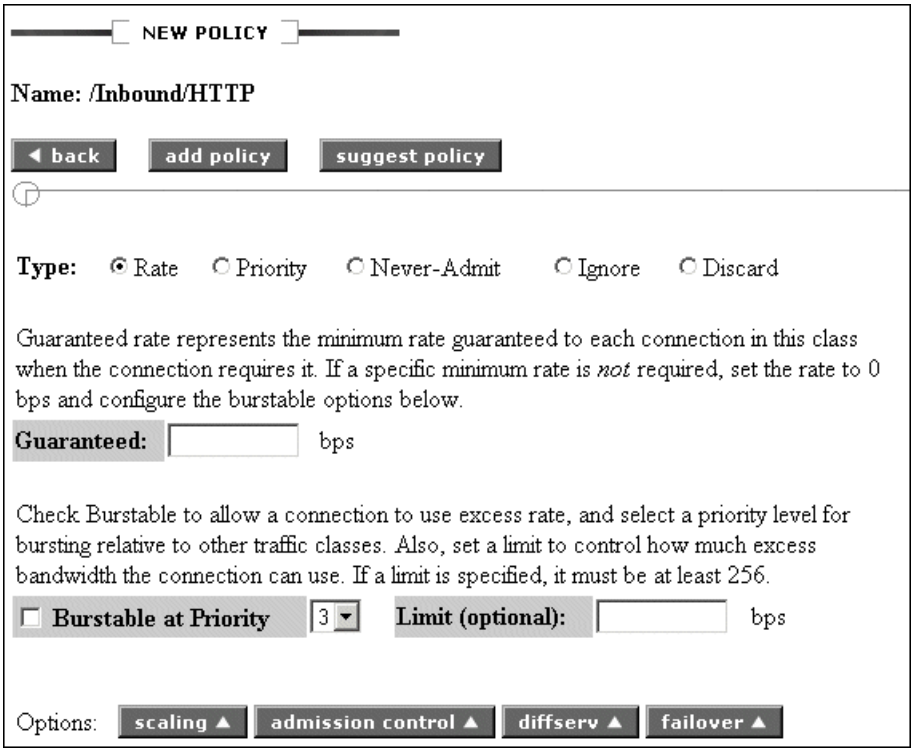

5. Configure the policy's attributes, as described in the following table:

Field	Description
Guaranteed	The minimum guaranteed rate given to each connection. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions). When a specific minimum rate is not required, set the rate to 0 bps and configure the **Burstable** option. The guaranteed rate must be a minimum of 1024 bps. **Tip:** You can use the **traffic flow** command in the CLI to see how many concurrent flows of a specific type are present. This will help to estimate guaranteed rate requirements. **Note:** If the traffic class has a partition, be sure that the guaranteed rate setting does not exceed the committed partition size. To determine the committed partition size, look at the **Size** field in the *Partition* window. If the policy's guaranteed rate is higher than the committed partition size, each flow will be "squeezed" to less than 1024 bps, thereby invoking admission control. For information on admission control, see Apply Admission Control.
Burstable at Priority	When this checkbox is selected, the traffic class can access unused bandwidth at the specified priority level (0-7, where 7 is the highest priority). The priority level defines how this class will compete with other classes for excess bandwidth. Priority 7 is very assertive, so save it for the most important traffic types. See Priority Guidelines.
Limit (*optional*)	Caps the excess rate available to the selected class. If a limit is specified, it must be greater than the policy's guaranteed rate.

Certain types of traffic offer specialized rate policy options. The *Policy* window displays buttons for the following functions if they are relevant to the selected traffic type:

- Scaling
- Admission Control
- UDP Delay Bound
- Diffserv
- VLAN

6. Click **add policy**.

> **Note:** In order for rate policies to take effect, traffic shaping must be enabled. See Enable/Disable Traffic Shaping.

Set a Priority Policy in the Advanced UI

Priority policies are used to establish a priority for traffic without specifying a particular rate. Use a priority policy for short, non-continuous traffic flows or traffic that does not burst. The priority determines the importance of the traffic class; you specify a value from 0 to 7, where 7 is the highest priority. The Inbound/Default and Outbound/Default classes are factory-set to a priority of 3.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To apply a priority policy to a traffic class:

1. Click the **Manage** tab.

2. In the left window pane, select the class to which you want to apply a rate policy.

3. Click **policy** and select **add**. The *New Policy* window appears.

4. Click **priority** in the *New Policy* window. **show screen...**

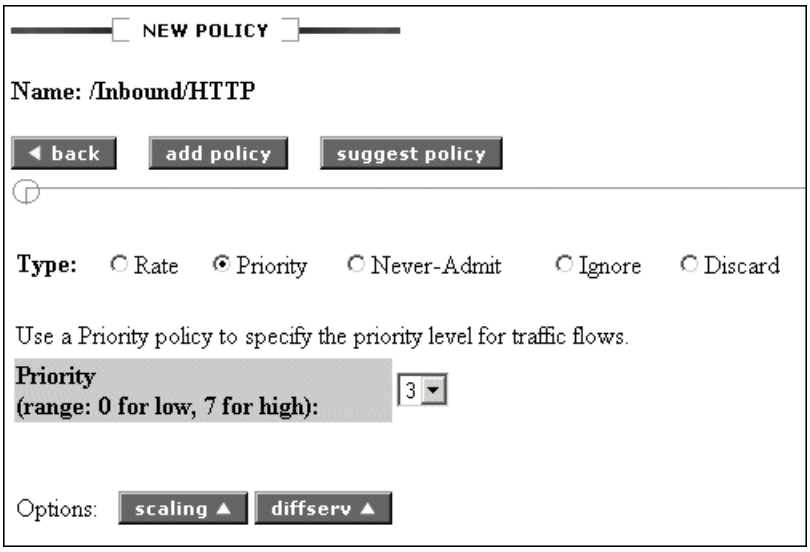

5. Select a priority (0 through 7) from the pull-down menu. Note that 0 is the lowest priority and 7 is the highest. See Priority Guidelines.

Certain types of traffic offer specialized priority policy options. The *Policy* window displays buttons for the following functions if they are relevant to the selected traffic type: Diffserv and Scaling.

6. Click **add policy**.

> **Note:** In order for priority policies to take effect, traffic shaping must be enabled. See Enable/Disable Traffic Shaping.

Priority Policy Examples

Use priority policies to prioritize specific applications that need preferential or low-priority treatment. The following table lists some typical uses of priority policies.

Traffic Class	Recommended Priority Level	Explanation
DNS LDAP	6	Directory lookups typically are very small, so let them through quickly to start the transaction.
SNMP	2	Network management flows usually do not burst and are not time sensitive.
TN3270 TN5250	7	These host-access applications have small, interactive, time-sensitive flows.

Set a Never-Admit Policy in the Advanced UI

A *never-admit policy* tells PacketShaper to enforce admission control at the beginning of each flow. While this policy type is named "never-admit," which implies that the connection is refused, there is an option in the case of web traffic classes only to redirect the user to another web page.

For TCP web traffic classes, use a never-admit policy to notify users that a website is unavailable or to redirect these users to a secondary site. For TCP non-web classes, use a never-admit policy to simply refuse a connection. For UDP traffic, use a discard policy instead.

> **Note:** When shaping is turned off, a class with a never-admit policy will display a rate of zero in the *Monitor Traffic* window. Packets are passed through the unit, but the rate values make it appear that the traffic is not passing through the unit.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To apply a never-admit policy to a traffic class:

1. Click the **Manage** tab.

2. In the left window pane, select the class to which you want to apply a never-admit policy.

3. Click **policy** and select **add**. The *New Policy* window appears.

4. Click **Never-Admit** in the *New Policy* window. **show screen...**

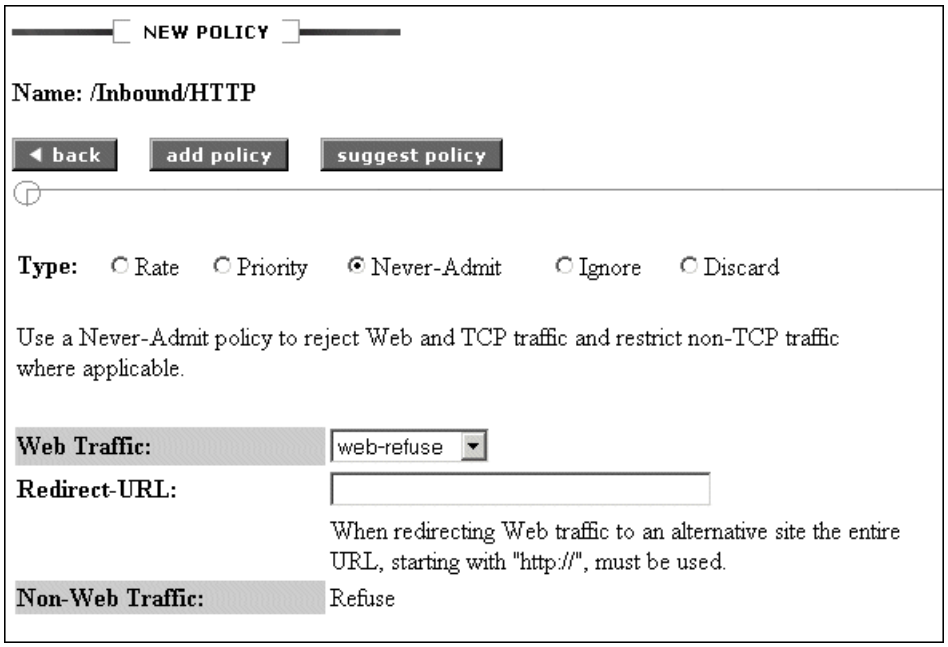

5. For web traffic, select one of the options below:

web-refuse to refuse the connection
-or-
web-redirect to redirect the user to another web page

6. If you selected **web-redirect**, specify the URL in the **Redirect-URL** text box. Specify the full URL for example, *http://www.mycompany.com/main.htm*.

If you omit the "http:" tag, the string will be appended to the end of the previously requested page, which is usually not the behavior you want.

7. Click **add policy**.

> **Note:** In order for never-admit policies to take effect, traffic shaping must be enabled. See Enable/Disable Traffic Shaping.

Special-Case Never-Admit Considerations

- A web never-admit policy can be applied to the *response* side of a flow. All other never-admit policies must be applied to classes on the *requesting* flow. Since admission control takes effect at the beginning of the connection, PacketShaper must react to the requesting flow.
- A never-admit policy cannot be applied to a non-IP protocol for example, AppleTalk. For these traffic types, use a discard policy.
- The procedure documented here applies to the case where the class being redirected is defined by a single web server. To redirect based on *client* IP address, not *server* IP, you'll need to create two classes: one that redirects the hosts (an HTTP class based on a host list of client IP addresses, with a never-admit policy that redirects to a website) and the

other class to classify traffic to the redirected website (HTTP class with IP address of the web server). The second class should be made an exception class.

Never-Admit Policy Example

A never-admit policy can be used to redirect certain web users to alternate URLs. For example, you might redirect a web request to a URL that presents a customized message.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To redirect web users to an alternate web page:

1. Click the **Manage** tab.

2. In the left window pane, select an HTTP traffic class.

> **Note:** The redirect option works only on the *response* side of an HTTP flow, not the request side. For example, specify the web-redirect option for a traffic class such as: */Inbound/HTTP/products.com*

3. Click **policy** in the *Traffic Class* window and select **add**.

4. Click **Never-Admit** in the *Policy* window. **show screen...**

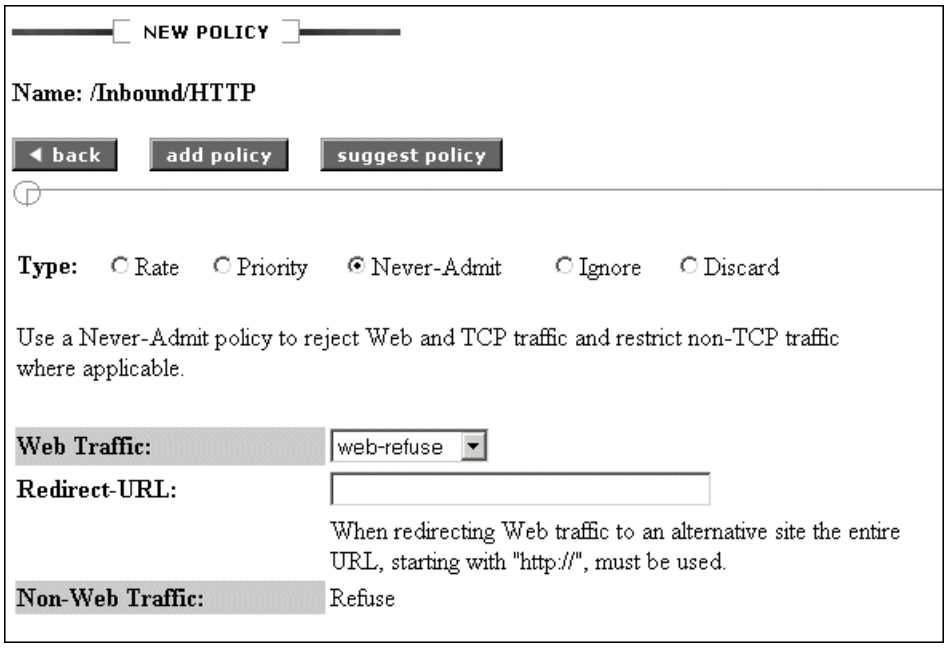

5. Select **web-redirect** from the **Web Traffic** drop-down list, then specify a redirect URL. Be sure to include the "http://" prefix in the URL specification.

6. Click **add policy**.

Set an Ignore Policy in the Advanced UI

Use an *ignore policy* to treat a traffic type as pass-through traffic so it will not be counted as part of the link traffic under management.

If PacketShaper sees traffic that is *not* destined for the router of the link that the unit is managing, this traffic should be ignored by PacketShaper, and should not be counted against the capacity of the link. For example, traffic to and from an Intranet server will never consume router bandwidth, so it should be ignored.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To apply an ignore policy to a traffic class:

1. Click the **Manage** tab.

2. In the left window pane, select the class to which you want to apply an ignore policy.

3. Click **policy** and select **add**.

4. Click **ignore** in the *New Policy* window. **show screen...**

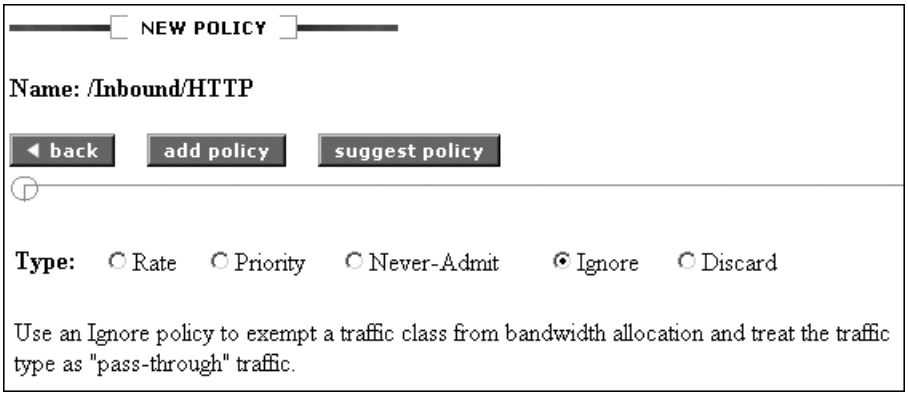

5. Click **add policy**.

> **Note:** Unlike other types of policies, ignore policies are in effect regardless of whether traffic shaping is enabled.

Set a Discard Policy in the Advanced UI

Use a *discard policy* to toss all packets for a traffic class. You may want to block a service or application because it is:

- Consuming too much of your network bandwidth
- Considered nonessential to your business
- Prohibited on your network

A discard policy is recommended for blocking UDP traffic classes because UDP traffic is not session oriented. For TCP or web traffic, you would typically use a never-admit policy to block traffic and inform the user that the session was disallowed. If you use a discard policy for TCP traffic, the discarded packets will cause TCP time-outs resulting in a long wait period before the user gets any feedback. In some cases this may be desirable. For example, you may want to "time out" someone attempting to crack your web server.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To apply a discard policy to a traffic class:

1. Click the **Manage** tab.
2. In the left window pane, select the class to which you want to apply a discard policy.
3. Click **policy** and select **add**.
4. Click **discard** in the *New Policy* window. **show screen...**

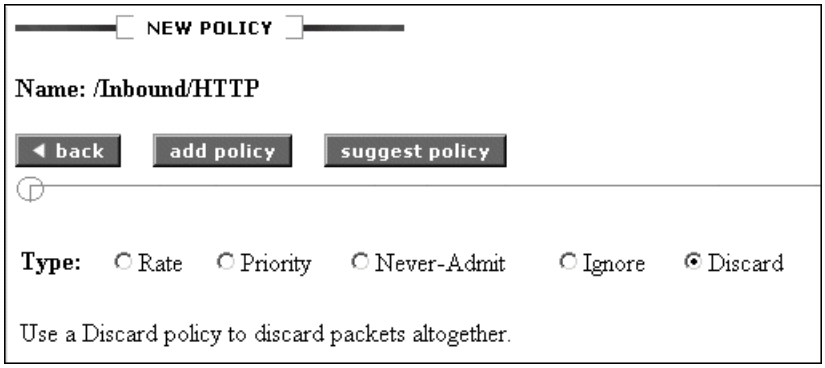

5. Click **add policy**.

> **Note:** In order for discard policies to take effect, traffic shaping must be enabled. See Enable/Disable Traffic Shaping.

Caution: Be careful when applying a discard policy, because it may have unexpected results. For example, if you configure the /Inbound/Default class to discard packets, PacketShaper will discard packets for all classes in the Inbound tree that do not have their own policies.

Modify a Policy

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To modify the settings of an existing policy:

1. Click the **Manage** tab. The *Traffic Class* window appears.
2. In the left window pane, select the class name.
3. Click **policy** in the *Traffic Class* window and select **open**. The *Policy* window appears.
4. Update the settings.
5. Click **apply changes** to update the policy.

Delete a Policy

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To delete a policy:

1. Click the **Manage** tab. The *Traffic Class* window appears.
2. In the left window pane, select the class name.
3. Click **policy** in the *Traffic Class* window and select **delete**.
4. Click **OK** in the confirmation box to remove the policy.

Apply Admission Control

The admission control mechanism for a rate policy determines what happens when there isn't enough bandwidth to satisfy guaranteed rate requests. Admission control lets you set a minimum level of service and even refuse connections when congestion occurs or a persistent backlog exists.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Class Tree** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To apply an admission control to a rate policy:

1. Open the rate policy you want to modify. (See Modify a Policy.)

2. Click **admission control** in the *Policy* window. The *Policy:Admission Control* window opens. **show screen...**

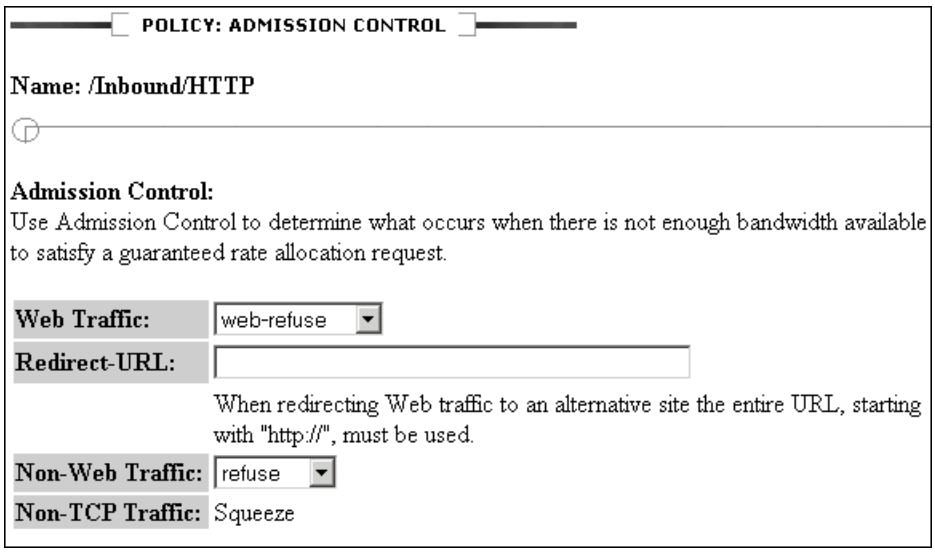

3. Select the type of control appropriate for the traffic class:

Traffic Type	Admission Control Options	Description
TCP Non-Web	refuse	Refuse new connection requests until enough bandwidth is available; connections are refused without user notification
	squeeze	Give new connections a minimal "trickle" rate of less than 1024 bps
	web-refuse	Refuse new connection requests until enough bandwidth is available; a refusal results in an HTTP error message
TCP Web	web-squeeze	Give new connections a minimal "trickle" rate of less than 1024 bps
	web-redirect	Redirect HTTP traffic to a user-specified URL; the redirect option is available only for HTTP classes on the *response* side of the HTTP flow, not the *request* side

Traffic Type	Admission Control Options	Description
Non-TCP	squeeze	Give new connections a minimal "trickle" rate of less than 1024 bps

4. If you selected **web-redirect**, specify the URL in the **Redirect-URL** text box. Specify the full URL for example, *http://www.mycompany.com/main.htm*.

If you omit the "http:" tag, the string will be appended to the end of the previously requested page, which is usually not the behavior you want.

5. Click **OK**.

6. Click **apply changes** in the *Policy* window.

Set Up UDP Latency Control

For UDP applications and non-IP protocols such as IPX and AppleTalk, you can use a delay bound in a rate policy to provide flow control. Since UDP does not rely on acknowledgments to signal successful receipt of data, there is no concept of a flow-control window. Packets are transmitted one after another without regard to whether or not the recipient can handle them. However, "well-behaved" UDP applications communicate packet loss to the sender so that it can slow down.

PacketShaper uses a *UDP latency control mechanism* to rate-control individual UDP traffic flows and minimize packet loss. PacketShaper accumulates incoming UDP packets on a flow-by-flow basis when they are not scheduled for immediate transfer. With the UDP latency control mechanism, you define a *delay bound* how long the packets can remain buffered before they become too old to be useful. If UDP flows don't get sent immediately (because of link congestion, for example), they are placed in a buffer or *queue*. UDP flows stay in the queue until they are sent or until the delay bound time is exceeded, in which case the packets are dropped.

> **Note:** Unless you have specific requirements for buffering non-TCP traffic, it is recommended that you do not change the delay bound value, as it has been optimized for most network environments.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To set latency control for a UDP class:

1. If you haven't done so already, create a rate policy for the UDP class.
2. Click **delay bound** in the *Policy* window. **show screen...**

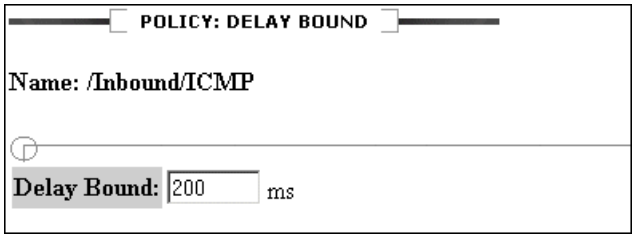

3. Define the maximum length of time that PacketShaper will retain UDP packets (1 to 10,000 milliseconds).
4. Click **OK**.
5. Click **apply changes** in the *Policy* window.

Set Up Rate Scaling

A rate policy can be configured to scale guaranteed and excess bandwidth allocation for low-speed and high-speed connections. PacketShaper detects the speed of an incoming flow and uses the policy's low-speed and high-speed rate settings to scale a minimum rate for the connection.

When PacketShaper examines a flow, it can discriminate between a client with a modem connection and a client with a higher-speed connection. The rate given to the client is based on a sliding scale, which is dependent on the client's speed. PacketShaper monitors the speed throughout the life of a connection and adjusts bandwidth allocation as the effective connection speed changes.

Take, for example, a class that has a rate policy with guaranteed rate settings of:
Low-speed = 10 Kbps; High-speed = 30 Kbps. In this example, a slower client would be guaranteed a rate value on the low end (perhaps 15 Kbps) and a faster client would be guaranteed a rate on the high end (perhaps 25 Kbps). Excess rate is distributed with a similar scaling factor, with slower clients receiving a smaller amount than high-speed clients.

This advanced tuning feature is rarely used, as PacketShaper automatically determines optimum bandwidth-allocation. However, if you have applications that require more precise control, use this feature to pace excess bandwidth allocation.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To set up rate scaling for a rate policy:

1. Open the rate policy you want to modify. (See Modify a Policy.)

2. Click **scaling** in the *Policy* window. The *Policy:Scaling* window opens. **show screen...**

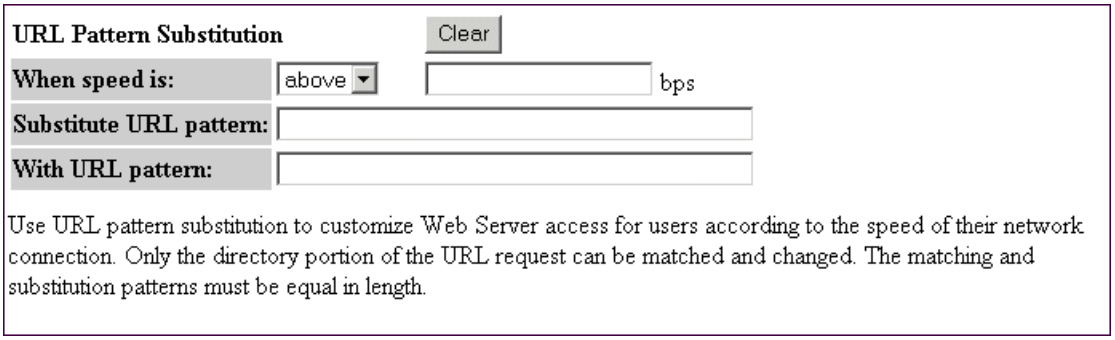

3. Choose **Explicit** for the Excess Rate to enable rate scaling.

4. In the **Low** text box, enter the desired rate for this class' low-speed connections. The minimum value allowed for the low-speed excess rate setting is 1024 bps. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions).

5. In the **High** text box, enter the desired rate for this class' high-speed connections.

6. In the **Limit** text box, enter the maximum excess rate that can be used by this class' low-speed and high-speed connections.

7. Click **OK**.

8. Click **apply changes** in the *Policy* window.

Set Up Web Redirection

This advanced policy option, also known as URL substitution, is available for *web traffic classes* only for example, HTTP or EntryPoint.

PacketShaper can detect the speed of a web connection at the first HTTP *get* request. You can configure PacketShaper to re-map the requested URL by substituting a URL that is more appropriate for the speed of the connection.

Using this feature, you can tailor the user's access experience and provide richer content to the users who can handle it. For example, if a high-speed user is accessing your website, PacketShaper can replace the plain and simple version of your home page banner with a multimedia banner. Conversely, when a low-speed user accesses your website, PacketShaper can provide basic, text-only content to reduce the download wait period.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

To set up web redirection for a web traffic class:

1. Open the rate or priority policy you want to modify. (See Modify a Policy.)

2. Click **scaling** in the *Policy* window. The *Policy:Scaling* window opens.

3. Scroll down to the **URL Pattern Substitution** section. **show screen...**

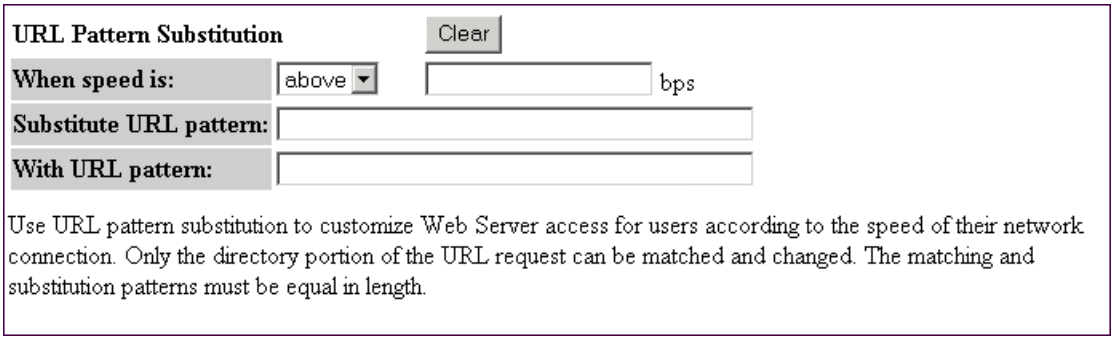

4. Specify the parameters as described in the following table.

Field	Description
When speed is	Specify **above** or **below** a connection speed to indicate when the URL should be substituted. For example, if the speed is below 9600, you might want to redirect the user to a text-only URL.
bps (speed)	The connection speed in bps that, in conjunction with **above** or **below**, triggers the content substitution
Substitute URL pattern	The current URL pattern, which will be substituted with a new pattern. Wildcard patterns are not supported. The maximum URL length is 128 characters. This URL string is compared with the pattern in the /path/filename portion of a URL. PacketShaper ignores the http://hostname portion of a URL when performing matching or substitution.
With URL pattern	The URL that you specify for substitution must be the same length that is, the same number of characters - as the original URL. The maximum URL length is 128 characters.

5. Click **OK**.

6. Click **apply changes** in the *Policy* window.

Substitute Diffserv Values

The Differentiated Services (Diffserv) standard outlines an approach for giving precedence to specific traffic types, providing a scalable solution to managing traffic. If you have classified traffic by Diffserv, you can change the precedence or service level of an application by creating policies to change the code point value or substitute different COS/TOS values.

For information on creating classes based on Diffserv values, see Classify Diffserv Code Point Traffic and Classify Diffserv COS/TOS Traffic. For Diffserv details, see RFC 2474 at http://ietf.org/rfc.html.

Diffserv is an option that is available in both rate and priority policies.

PolicyCenter To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

Substituting a Diffserv Code Point Value

To substitute a different code point value:

1. Open the priority or rate policy. (See Modify a Policy.)
2. Click **diffserv** in the *Policy* window. The *Policy:DiffServ* window opens. **show screen...**

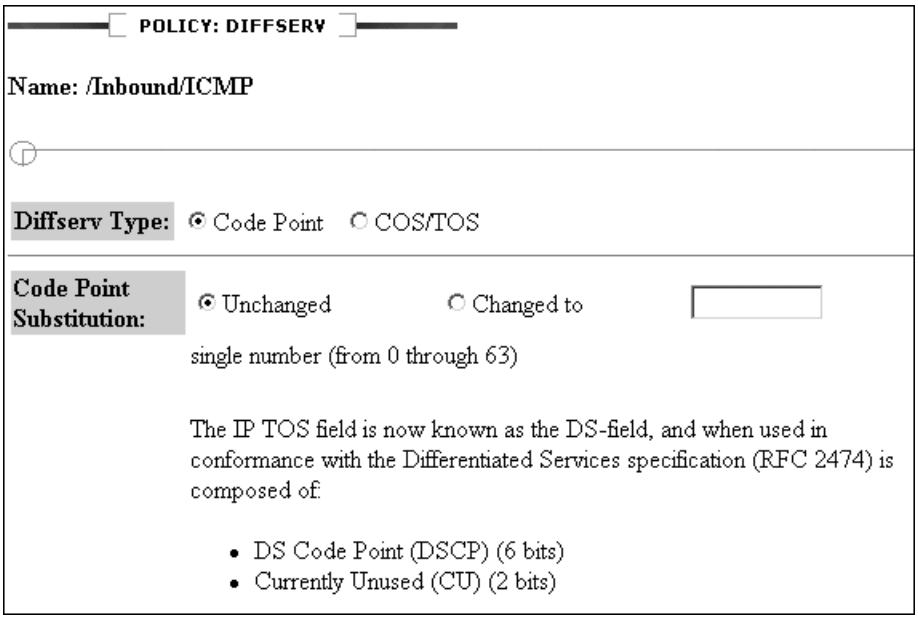

3. Make sure **Code Point** is selected for the Diffserv Type.
4. Select **Changed to** and enter a value (0-63) in the text box.
5. Click **OK**.
6. Click **apply changes** in the *Policy* window.

Substituting a Diffserv COS/TOS Value

The COS/TOS Diffserv option lets you override a packet's Type of Service (TOS) value before forwarding an IP packet, thus controlling the priority when it reaches the router.

To override a packet's precedence before forwarding the packet, thus controlling the priority when it reaches the router, use IP precedence class-of-service (COS) substitution and specify a preference level.

Using TOS substitution, you can reassign the service level for an application, based on your business perspective. For example, interactive applications with small data packets (for example, Telnet) usually require a minimum delay, while file transfer applications like FTP require maximum throughput.

To specify a COS/TOS substitution:

1. Open the priority or rate policy. (See Modify a Policy.)

2. Click **diffserv** in the *Policy* window. The *Policy:DiffServ* window opens.

2. Choose **COS/TOS** for the Diffserv Type. **show screen...**

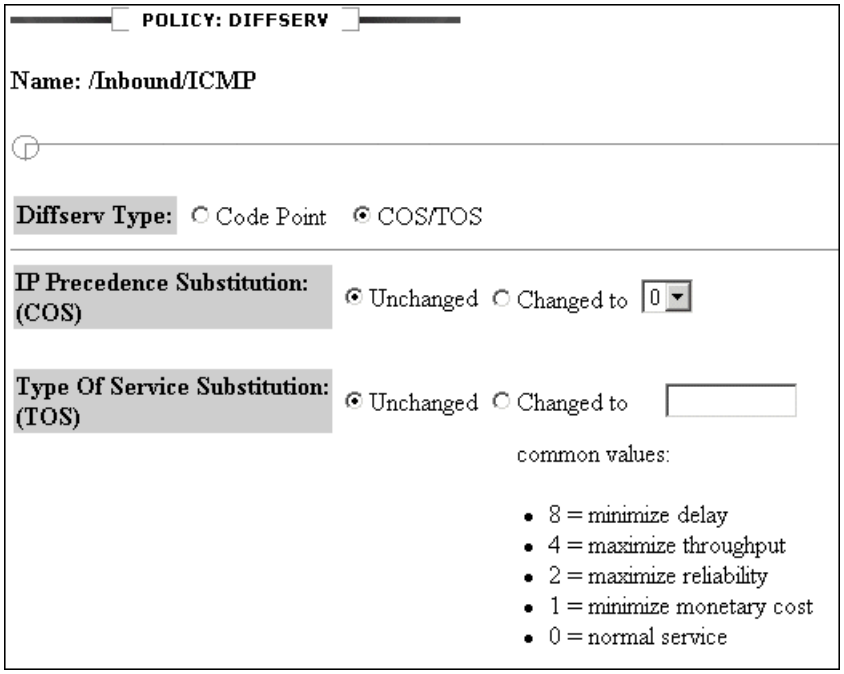

3. For IP Precedence Substitution, select **Changed to** and select an IP precedence (COS) value from 0 to 7, where 7 is the highest precedence.

4. For Type of Service Substitution, select **Changed to** and enter a TOS value.

5. Click **OK**.

6. Click **apply changes** in the *Policy* window.

Create VLAN Policies

You can create a policy to change the VLAN (virtual LAN) identification (802.1Q) or priority value (802.1p) on a packet. VLAN policies can be applied only to classes with rate or priority policies.

PolicyCenter To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window.

Change the VLAN Priority (802.1p)

To change the priority tag on an 802.1p class:

1. Open the priority or rate policy. (See Modify a Policy.)

2. Click **VLAN** in the *Policy* window. The *Policy:VLAN* window opens. **show screen...**

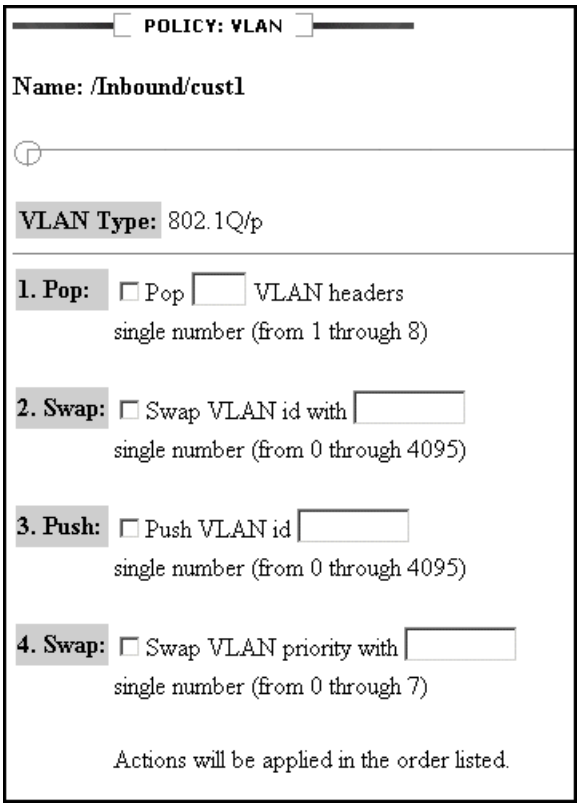

3. Select the **Swap VLAN priority** checkbox.

4. Enter the priority value (0 7).

5. Click **OK**.

6. Click **apply changes** in the *Policy* window.

PacketShaper will swap the topmost priority level on the VLAN stack with the specified priority value.

Change the VLAN Identification (802.1Q)

There are three operations you can perform on VLAN identification:

- *pop* removes the topmost label of the VLAN stack
- *swap* replaces the topmost ID of the VLAN stack with a different ID
- *push* puts an ID entry in a packet (and creates the stack if it doesn't exist)

A class can have a combination of push, swap, and pop operations in its VLAN policy. If more than one operation type is specified for a given class, they are executed in the following order: pop, swap, push. For example, a class might have a policy that specifies a swap, three pops, and a push. In this case, the three pops occur first, then the swap, and then the push.

To modify the VLAN identification tag on an 802.1Q class:

1. Open the priority or rate policy. (See Modify a Policy.)

2. Click **VLAN** in the *Policy* window. The *Policy:VLAN* window opens. **show screen...**

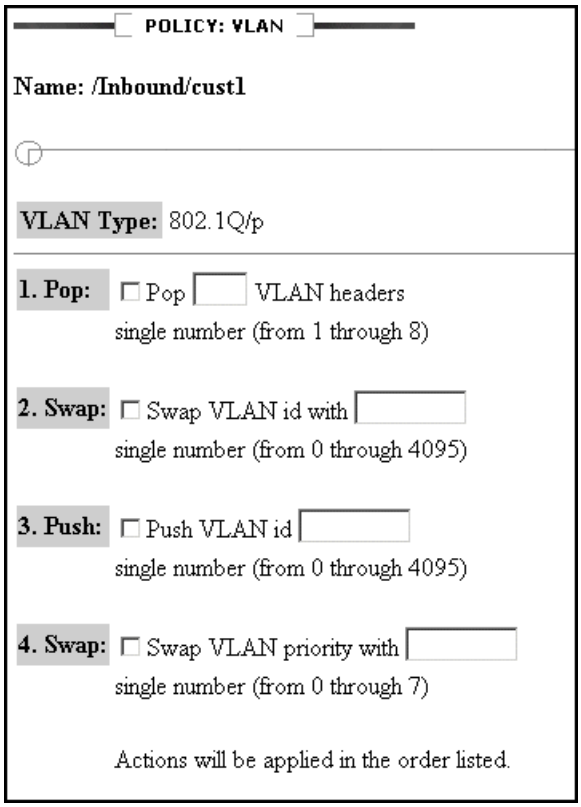

3. To pop off the topmost label of the VLAN stack, select the **Pop** checkbox and enter the number of times (1 8) you want to pop off the label.

4. To swap the topmost ID of the VLAN stack with another ID, select the **Swap VLAN id** checkbox and enter the ID value (0 4095).

> **Note:** A VLAN ID swap policy will automatically zero out the existing VLAN priority. To keep an existing non-zero priority value or to set a priority, be sure to specify a VLAN priority swap policy as well.

5. To put an ID entry in a packet, select the **Push VLAN id** checkbox and enter the ID value (0 4095).

6. Click **OK**.

7. Click **apply changes** in the *Policy* window.

Enable/Disable Traffic Shaping

When shaping is on, traffic is classified and measured, and policies and partitions actively control bandwidth allocation. When shaping is off, traffic is classified and measured, but it is not actively managed.

> **Note:** Partitions and policies (with the exception of ignore policies) do not take effect until traffic shaping is enabled.

To perform this task from PolicyCenter, you must first select a unit or draft configuration draft in the **Configurations** window. Then select the **Setup** tab from the right pane of this window, and proceed to step 2 of the procedure below.

To turn traffic shaping on or off:

1. Click the **Setup** tab. The *Basic Settings* screen is automatically displayed in the *Setup* window. **show screen...**

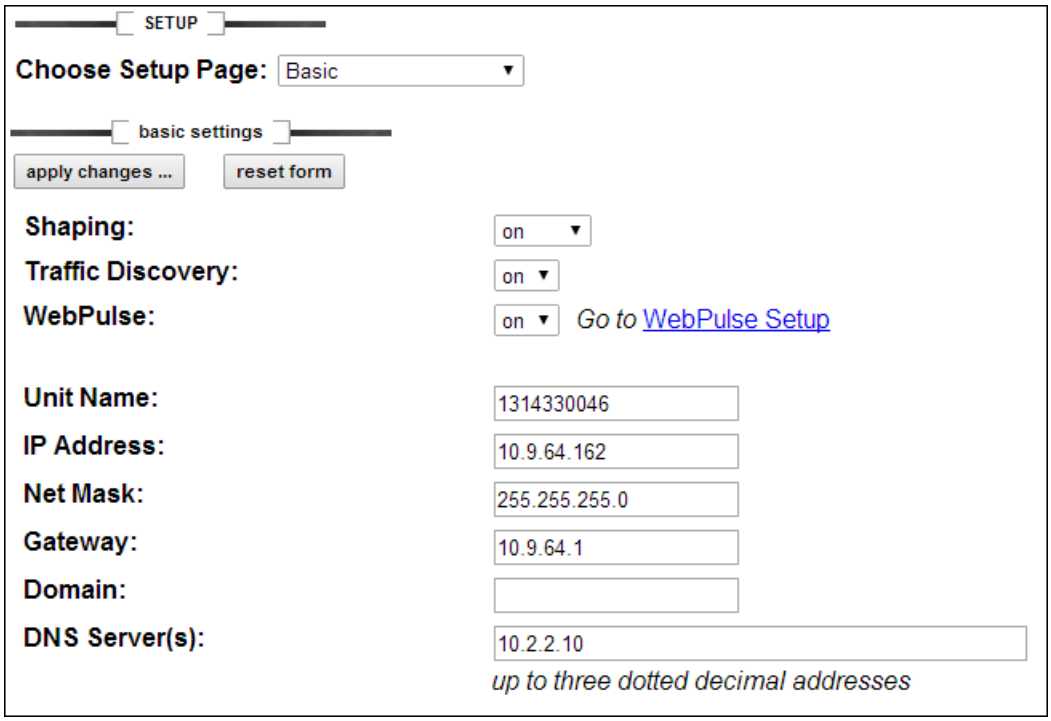

2. In the **Shaping** field, select **on** to enable shaping or **off** to disable.

3. Click **apply changes** to update the settings.

Graph and Report Tasks

Using PacketShaper s extensive reporting capabilities, you can view graphs of link and application usage and performance. PacketShaper collects data for over 120 different measurement variables; this data can be graphed in a variety of ways. For a list of all graphs available in PacketShaper, see Preconfigured Graphs.

This section includes tasks for creating a detailed traffic assessment report, graphing a single class, creating a multiple-graph report on a specific class, showing link utilization, and graphing multiple classes on a single chart.

Create a Traffic Assessment Report

Generate year-to-date, month-to-date, week-to-date, and daily traffic assessment reports of your Internet link. This detailed report can be viewed in a web browser, and downloaded and saved as a PDF. Use this report to find out the top applications being used on your network, which hosts and host pairs have the highest bandwidth usage, how much of your link is consumed by social networking traffic, and details about encrypted traffic.

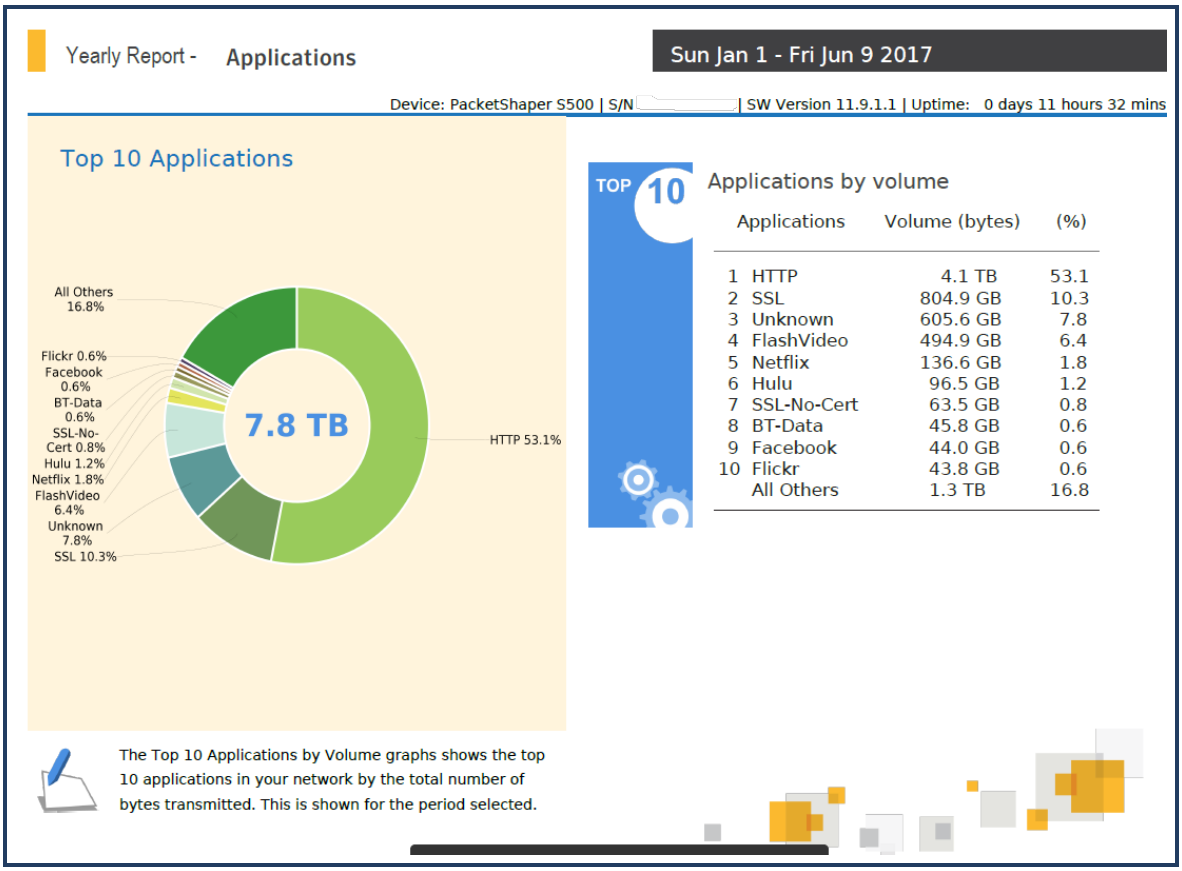

Data for all current reports is automatically updated every 15 minutes, except for periods when the PacketShaper is offline. When PacketShaper auto-generates the current set of reports, it overwrites existing reports with the latest data. For example, PacketShaper stores one daily report for each day; this daily report gets replaced with current data every 15 minutes throughout the day. If you want to keep a permanent record of a point in time, you can view today's daily report and save the PDF to your local drive, making sure to add the time of day to the filename.

Historical reports are stored as PDFs on the appliance and are available for viewing for five years. Report names include the date so that they are easily identifiable (for example, 201706 for the June 2017 monthly report and 20170608 for the June 8, 2017 daily report).

Included Reports

The reports included in the assessment are:

Report	Description	What to look for

Top 10 Applications

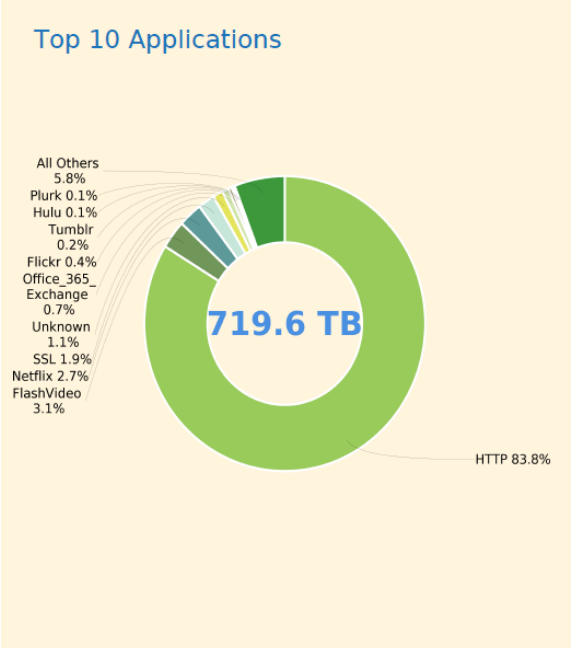

Applications by volume

	Applications	Volume (bytes)	(%)
1	HTTP	602.8 TB	83.8
2	FlashVideo	22.4 TB	3.1
3	Netflix	19.5 TB	2.7
4	SSL	13.8 TB	1.9
5	Unknown	8.3 TB	1.1
6	Office_365_Exc	5.4 TB	0.7
7	Flickr	2.7 TB	0.4
8	Tumblr	1.1 TB	0.2
9	Hulu	1.0 TB	0.1
10	Plurk	759.2 GB	0.1
	All Others	41.8 TB	5.8

A doughnut graph and list showing the 10 applications with the most traffic during the period.

- Is bandwidth consumption aligned with business objectives?
- Are recreational applications consuming too much bandwidth?
- Has PacketShaper identified any unsanctioned or risky apps?

The Top 10 Applications by Volume graphs shows the top 10 applications in your network by the total number of bytes transmitted. This is shown for the period selected.

Applications

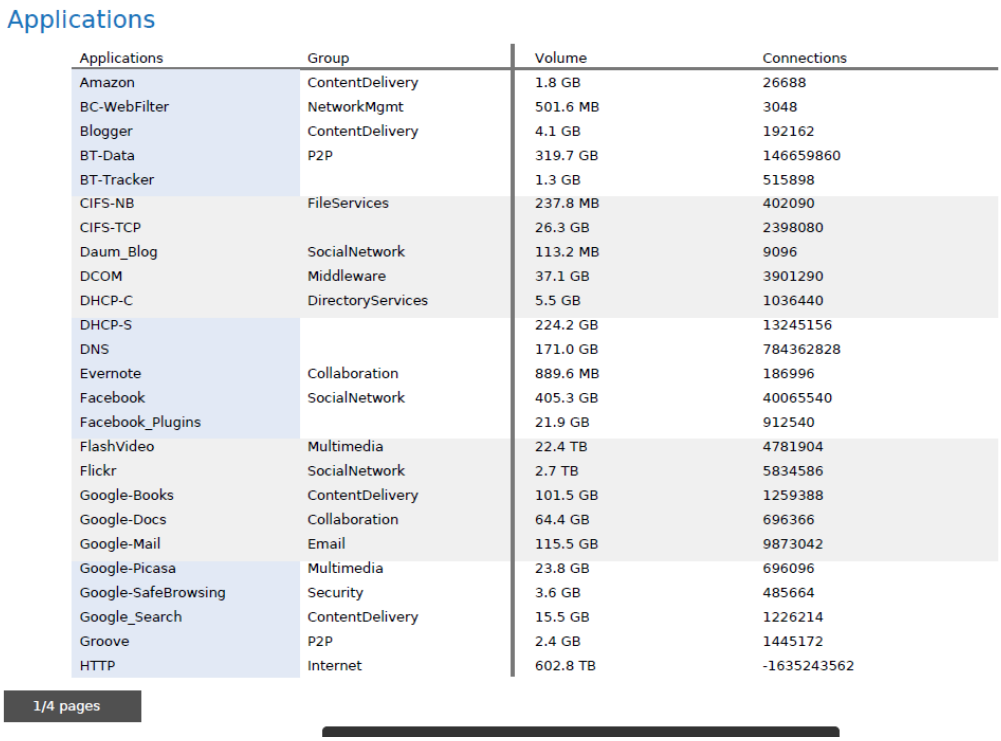

Applications

Applications	Group	Volume	Connections
Amazon	ContentDelivery	1.8 GB	26688
BC-WebFilter	NetworkMgmt	501.6 MB	3048
Blogger	ContentDelivery	4.1 GB	192162
BT-Data	P2P	319.7 GB	146659860
BT-Tracker		1.3 GB	515898
CIFS-NB	FileServices	237.8 MB	402090
CIFS-TCP		26.3 GB	2398080
Daum_Blog	SocialNetwork	113.2 MB	9096
DCOM	Middleware	37.1 GB	3901290
DHCP-C	DirectoryServices	5.5 GB	1036440
DHCP-S		224.2 GB	13245156
DNS		171.0 GB	784362828
Evernote	Collaboration	889.6 MB	186996
Facebook	SocialNetwork	405.3 GB	40065540
Facebook_Plugins		21.9 GB	912540
FlashVideo	Multimedia	22.4 TB	4781904
Flickr	SocialNetwork	2.7 TB	5834586
Google-Books	ContentDelivery	101.5 GB	1259388
Google-Docs	Collaboration	64.4 GB	696366
Google-Mail	Email	115.5 GB	9873042
Google-Picasa	Multimedia	23.8 GB	696096
Google-SafeBrowsing	Security	3.6 GB	485664
Google_Search	ContentDelivery	15.5 GB	1226214
Groove	P2P	2.4 GB	1445172
HTTP	Internet	602.8 TB	-1635243562

1/4 pages

An alphabetical list of the 100 applications (native as well as web apps) with the most activity on the network during the period. Shows the application group, volume of traffic during the period, and the number of connections.

- Are there any unwanted applications on the network? Set policies to block this traffic.

Report	Description	What to look for

Top 10 Host Pairs

	Local IP	Remote IP	Volume
1	14.100.9.254	172.20.9.254	299.2 GB
2	14.100.7.136	172.20.7.136	298.6 GB
3	14.100.8.136	172.20.8.136	298.3 GB
4	14.100.21.32	172.20.21.32	298.1 GB
5	14.100.3.15	172.20.3.15	298.0 GB
6	14.100.25.149	172.20.25.149	297.9 GB
7	14.100.27.16	172.20.27.16	297.5 GB
8	14.100.4.68	172.20.4.68	297.4 GB
9	14.100.23.173	172.20.23.173	297.3 GB
10	14.100.29.156	172.20.29.156	297.2 GB

The Top 10 inbound host pairs table displays the bandwidth utilization of the top 10 inbound pairs of hosts for the period selected.

A list of host pairs (local IP address and remote IP address) with the most bandwidth usage for the period. There are separate reports for inbound and outbound.

- What IPs are sending or receiving large amounts of traffic to/from a specific host/user?
- Is there a host pair with abnormally high usage? This could indicate a DDOS attack or data exfiltration.

Top 10 Hosts

	IP Address	Volume
1	2.60.13.23	447.3 GB
2	2.60.30.5	390.7 GB
3	2.60.17.15	348.5 GB
4	2.60.12.192	313.2 GB
5	2.60.19.134	311.4 GB
6	172.20.9.254	299.2 GB
7	172.20.7.136	298.6 GB
8	172.20.8.136	298.3 GB
9	172.20.21.32	298.1 GB
10	172.20.3.15	298.0 GB

The Top 10 sending hosts table displays the bandwidth utilization of the top 10 talkers during the period selected.

A list of host IP addresses that had the most bandwidth usage during the period. There are separate reports for sending hosts (talkers) and receiving hosts (listeners).

- What IPs (hosts/users) are generating the most traffic?
- Are traffic volumes "normal" or is there potential network abuse?
- Could high traffic volumes indicate a malware outbreak?

Report	Description	What to look for
Social Networking 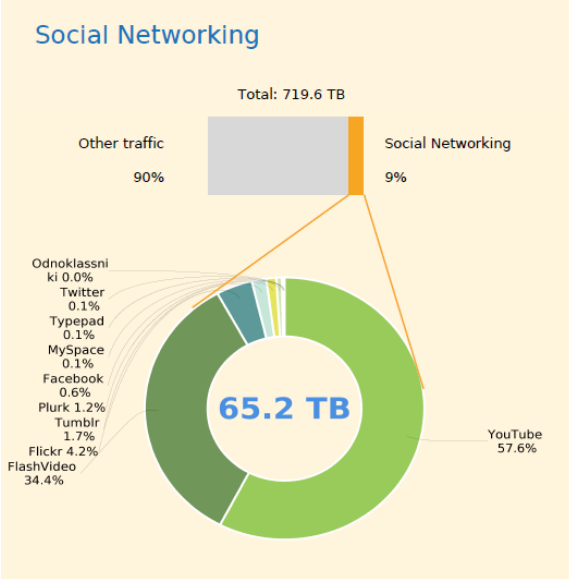 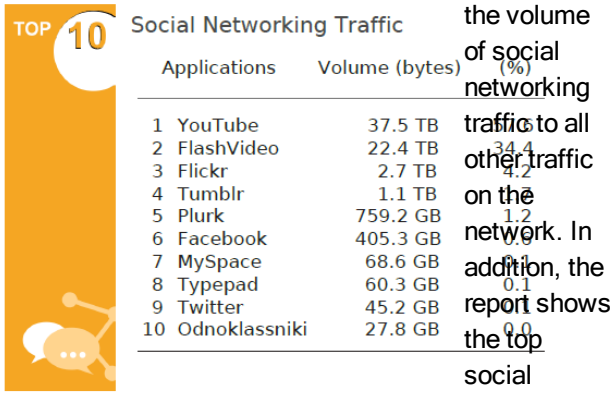 Top 10 social networking classified traffic	A doughnut graph that compares the volume of social networking traffic to all other traffic on the network. In addition, the report shows the top social networking applications that were used during the period.	■ How much bandwidth is consumed by recreational traffic? ■ Is excessive social networking traffic preventing business-critical traffic from getting appropriate priority and bandwidth? ■ Do appropriate use/compliance policies need to be revisited?
Encrypted Traffic 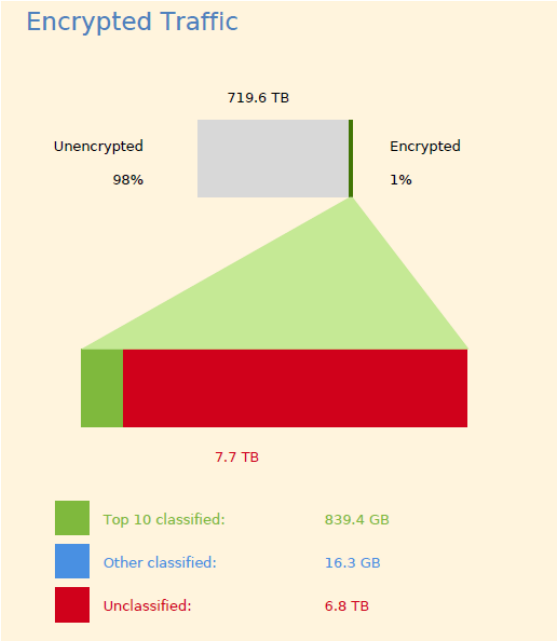 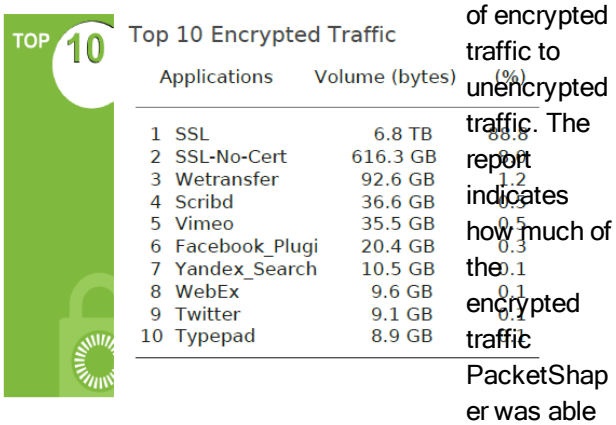 Top 10 encrypted and classified traffic	A graph that compares the volume of encrypted traffic to unencrypted traffic. The report indicates how much of the encrypted traffic PacketShaper was able to classify as well as a list of the top 10 classes of encrypted traffic.	■ Understand the "blind spot" represented by encrypted traffic. ■ Identify "SSL-No-Cert" traffic that represents a security risk. Note that incorporating a Symantec SSL Visibility appliance into your network can offer additional insight into encrypted traffic.

View a Traffic Assessment Report

1. In the Advanced UI, click the **Report** tab.
2. In the On-Box Reporting section on the left, click **Get Reports**.

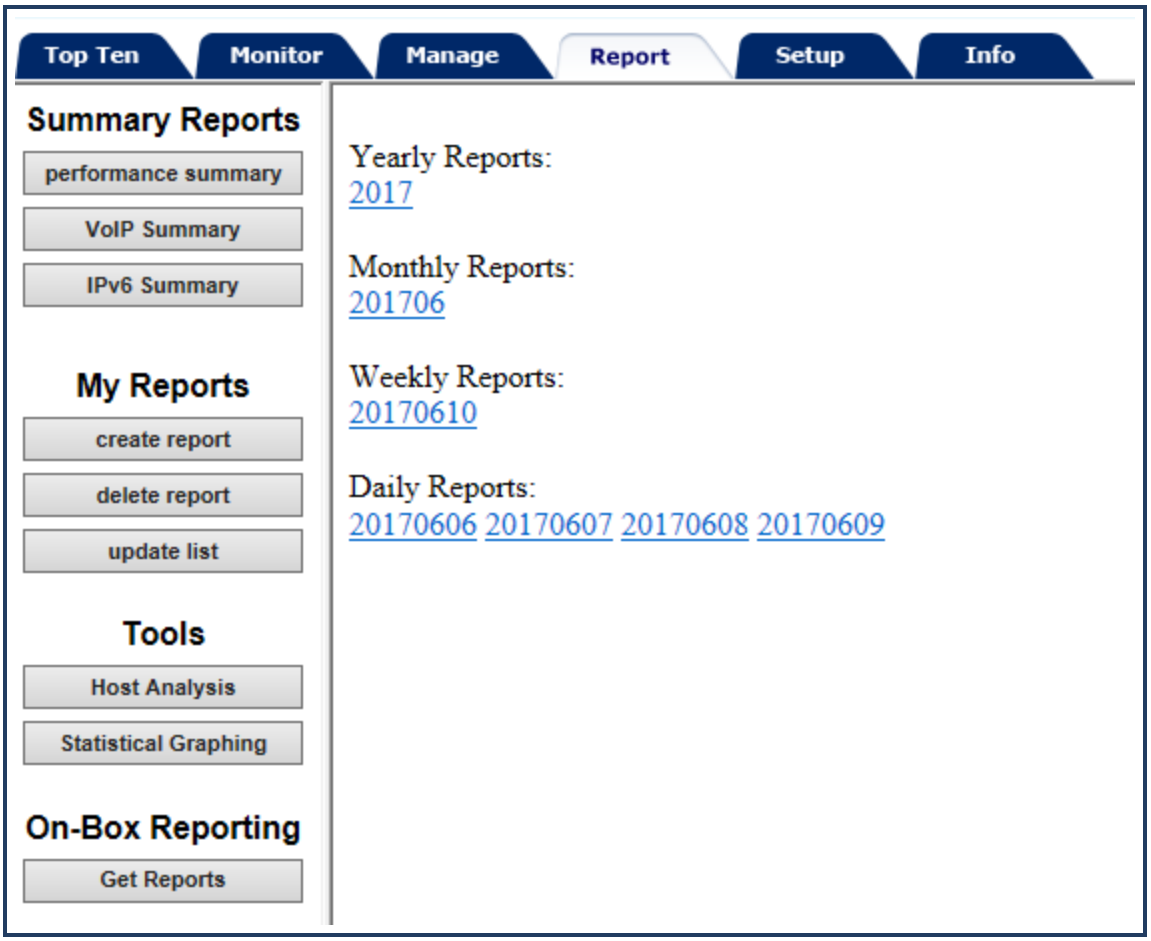

3. Select the report corresponding to the period you want to assess:
   - Yearly Reports (for example, click *2017*)
   - Monthly Reports (for example, click *201706* to view the June 2017 report)
   - Weekly Reports (for example, click *20170610* to view the report for the week ending June 10, 2017). Note that weekly reports begin on Sundays.
   - Daily Reports (for example, click *20170607* to view the daily report for June 7, 2017.

   A PDF of the report opens in a new tab or window.
4. (optional) Download or print the report.

Additional Information

- Report PDF files are stored in the /public/reports folder on the appliance for five years. You can use the PacketShaper File Browser to zip, download, or delete the files.

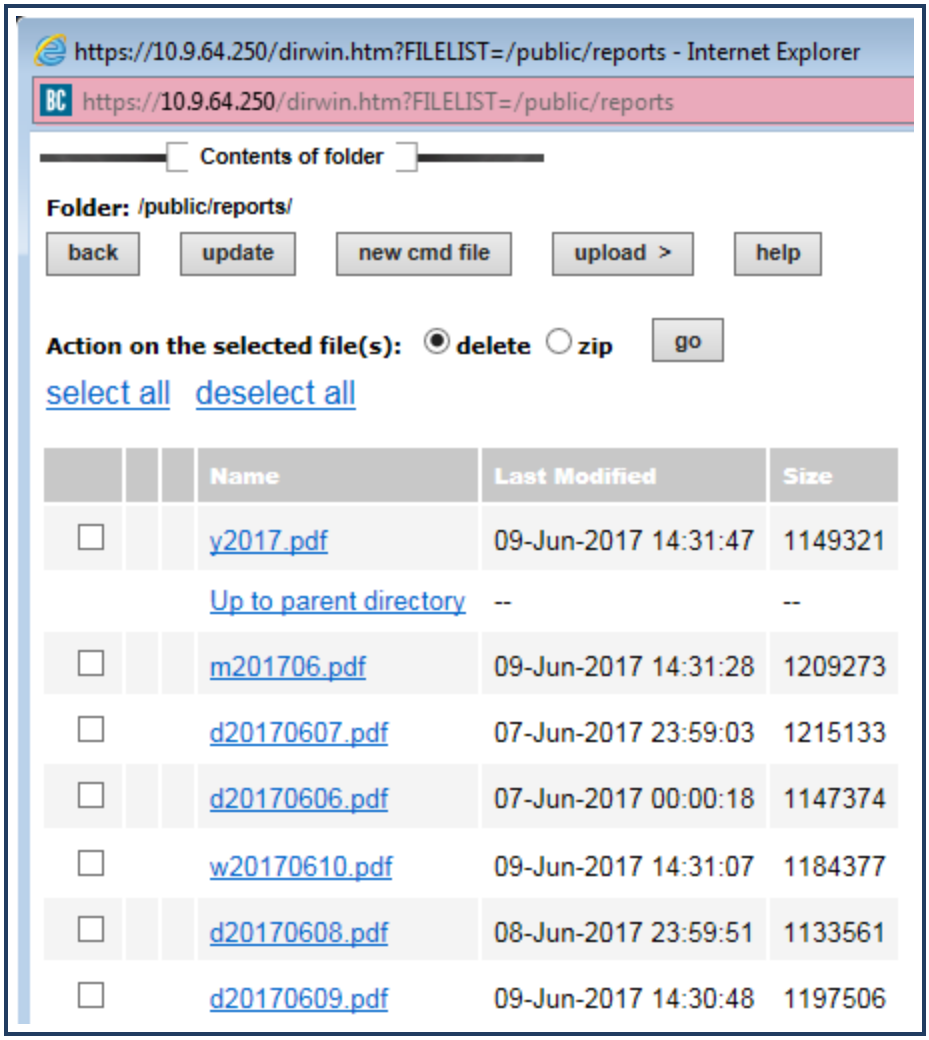

- If you downgrade PacketShaper to a pre-11.9.1 version, all PDF files of reports generated with the on-box reporting feature will remain on the hard drive and can be viewed or downloaded using the File Browser.

View the Top Ten Pie Chart

The Network Performance Summary report includes pie charts of the ten Inbound and ten Outbound classes that have generated the most traffic. The chart shows the relative portions of bandwidth allocated to the ten most active classes on the network and displays each class' average bandwidth usage in bits per second and its percentage of the total bandwidth on the link.

To view the top ten classes on your network:

1. Click the **Report** tab. The Network Performance Summary report displays. **show screen...**

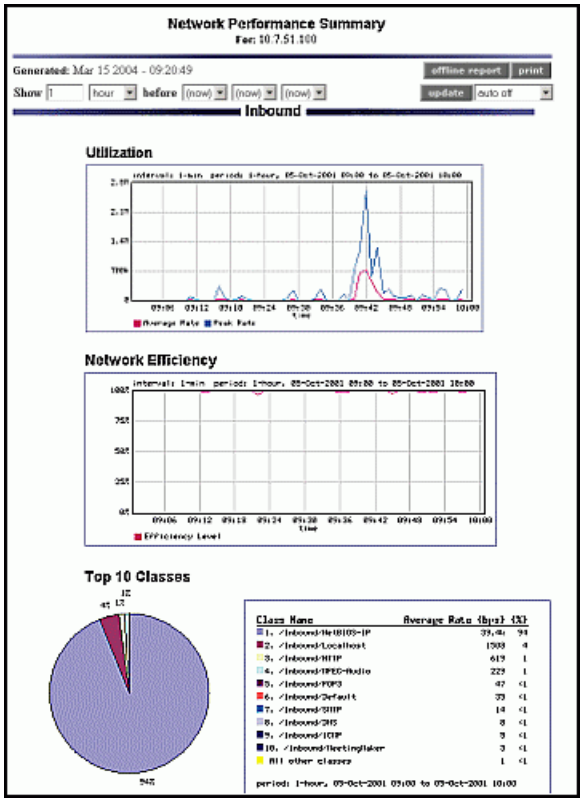

2. Scroll down until you see the Top Ten Classes pie chart (there's one for Inbound and another for Outbound).

The class names in the pie chart legend are hot links to the *Statistics:Report* page, offering a quick way to drill down and get more detailed information about a top-ten class. Note that this capability is disabled in PacketShapers with more than 2000 classes, in order to speed up the display of the Network Performance Summary.

3. To create a report on one of the classes in the pie chart, click the class name in the pie chart legend. **show screen...**

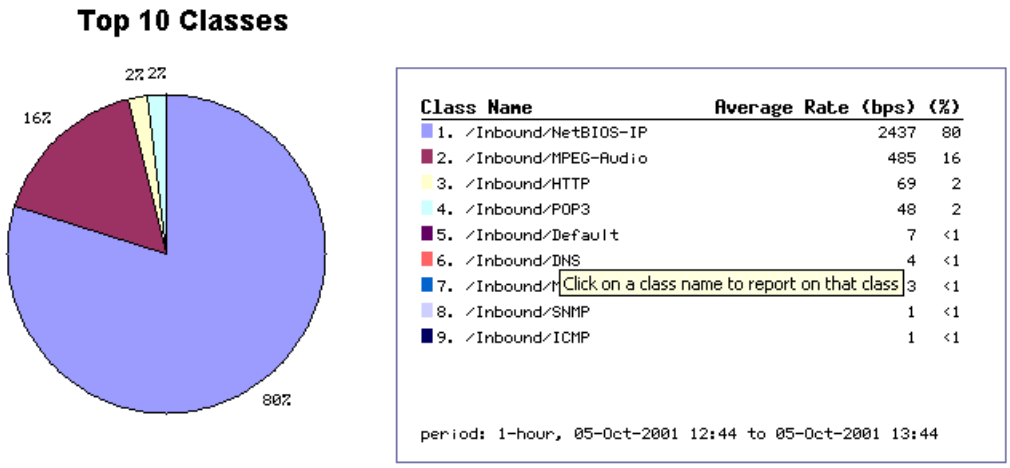

The *Statistics:Reports* screen appears. **show screen...**

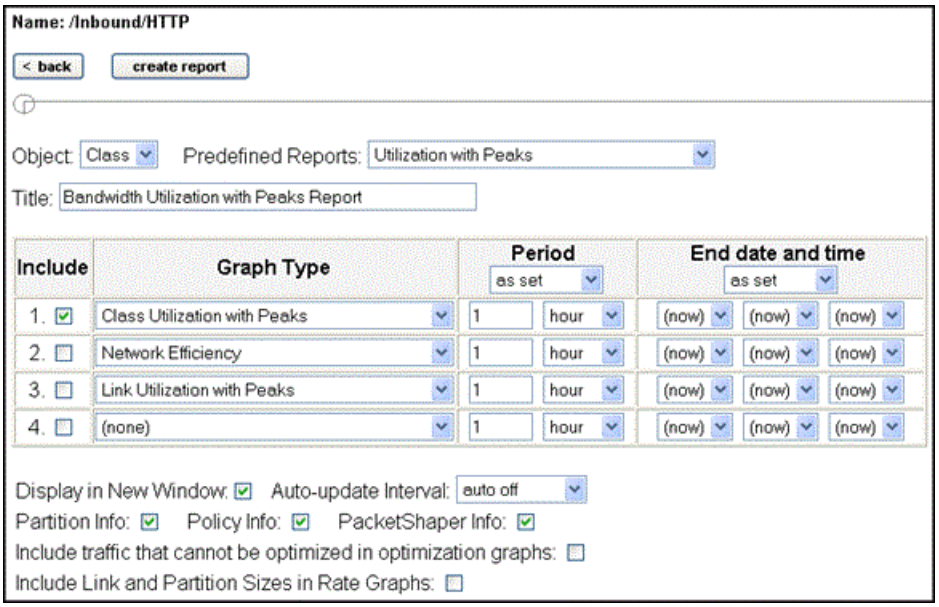

4. Change the period and/or end date, and then click **create report**. The Bandwidth Utilization with Peaks report for the selected class appears in a new window.

List the Top Ten Classes

To display a list of the Inbound and Outbound classes that have generated the most traffic, click the **Top Ten** tab. The default report, *Average Rate*, appears. **show screen...**

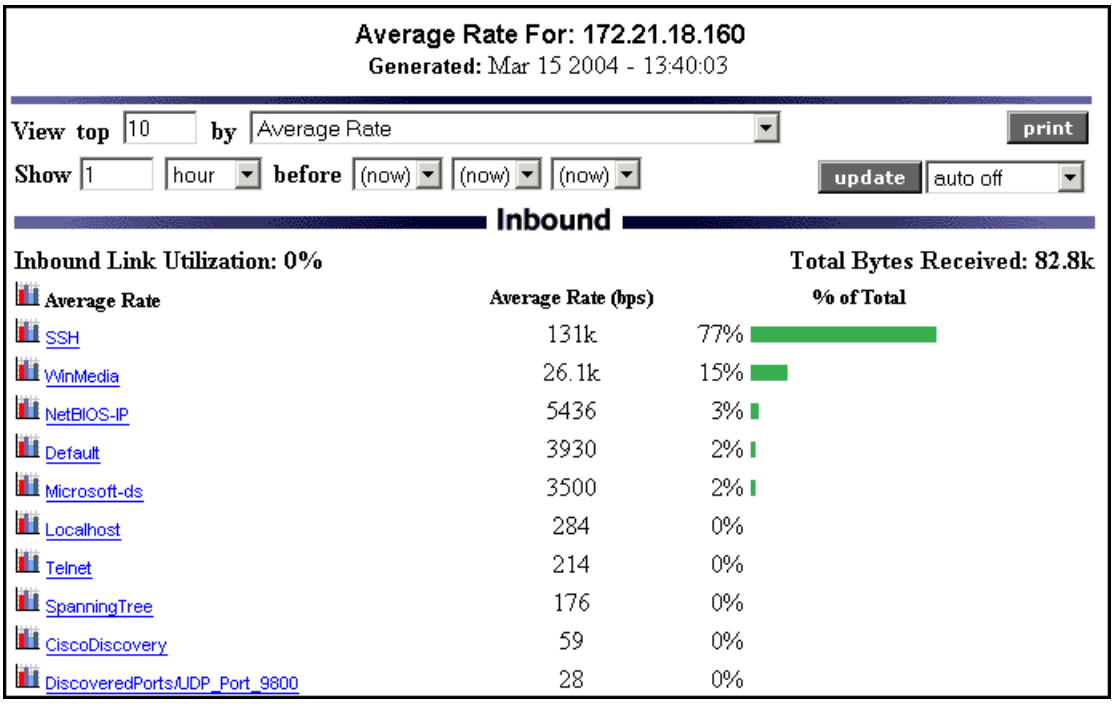

This report lists the following information for each top-ten traffic class:

Column	Description
	A link that displays a graphical report of the class' average and peak rates and network efficiency.
	The at the top displays a stacked utilization graph of the traffic classes displayed on the list. For example, if the top 15 classes are displayed, the stacked utilization graph will graph 15 classes.
Average Rate (bps)	The average bandwidth consumption in bits per second
% of Total	The percentage of total bandwidth consumed

The report displays the ten classes that have generated the most traffic in the last hour. To display a different number of classes or view a different report, see Select a Different View. To specify a different time period, see Select the Time Period.

In addition to the top classes, this report provides you with statistical data about your link utilization. **Inbound Link Utilization** indicates the percentage of your inbound link that was used in the time period and **Total Bytes Received** indicates the number of bytes that were received in the Inbound direction during the time period. These same types of statistics are provided in the Outbound direction as well.

Create a Graph

PacketShaper includes dozens of graphs which you can display for a class, partition, or link. You can graph bandwidth usage, network efficiency, top ten classes, and TCP connections. For descriptions and examples of all the graphs, see Preconfigured Graphs.

For your convenience, PacketShaper lets you create a graph from any of the following pages:

- Top Ten
- Monitor Traffic
- Manage

Creating Graphs from the Top Ten Page

From the Top Ten page, you have single-click access to a graphical report of a class' bandwidth utilization.

To view a report containing utilization graphs (Average and Peak Rates and Network Efficiency) for a class:

1. Click the **Top Ten** tab. The default report, *Average Rate*, appears. **show screen...**

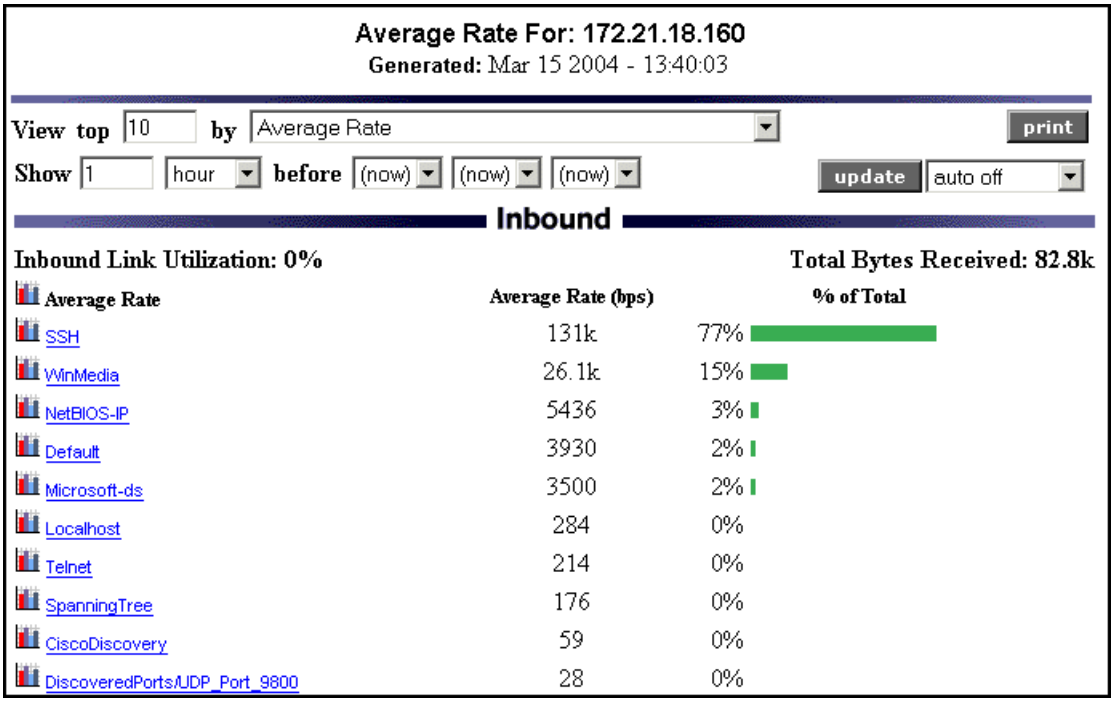

2. If the class doesn't appear on the Top Ten list, clear out the **View top** field and click **update**. All classes will be displayed.

3. Click the icon next to the class you want to graph.

   A window appears with a report for the selected class. This report contains two graphs: Average and Peak Rates and Network Efficiency. It uses the time period specified on the Top Ten page.

> **Note:** The at the top of the list (next to *Top Ten - Average Rate*) displays a stacked utilization graph of the displayed traffic classes.

Creating Graphs from the Monitor Traffic Page

The *Monitor Traffic* page offers a shortcut to the *Statistics:Reports* page so that you can quickly display a default graph (such as Class Utilization with Peaks for the last hour), or you can modify the report (select a different graph, change the period, etc.).

1. Click the **Monitor** tab.
2. Click the icon next to the class you want to graph. The *Statistics:Reports* screen appears. **show screen...**

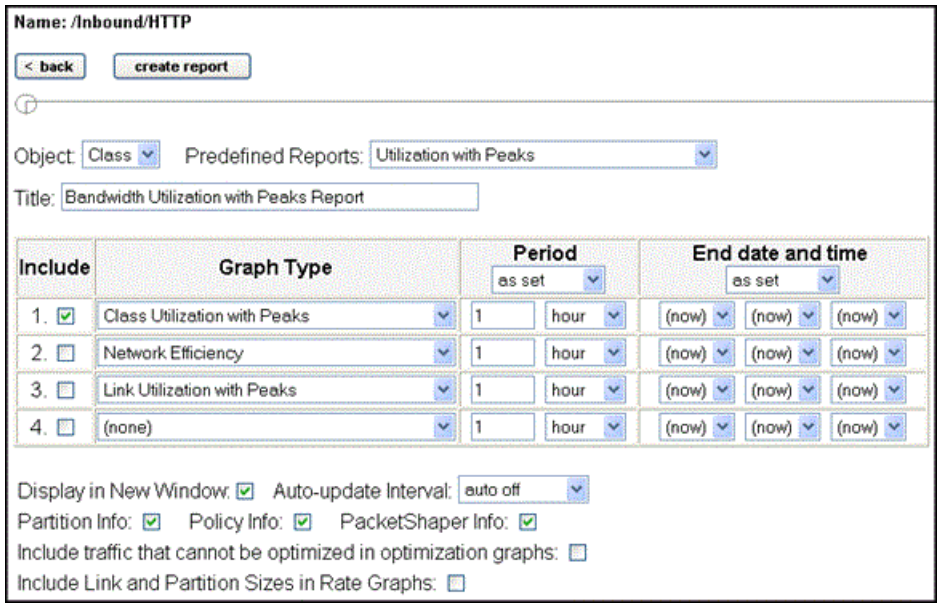

3. Click **create report** to display the default report.

   or

   Modify the report settings. See Creating a Graph from the Manage Tab, starting with step 5.

Creating Graphs from the Manage Tab

To create a graph of a specific link, partition, or class:

1. Click the **Manage** tab.

2. In the left window pane, select the link, partition, or class you want to graph.

3. Click **statistics** and choose **graph**. The *Statistics:Reports* screen appears. **show screen...**

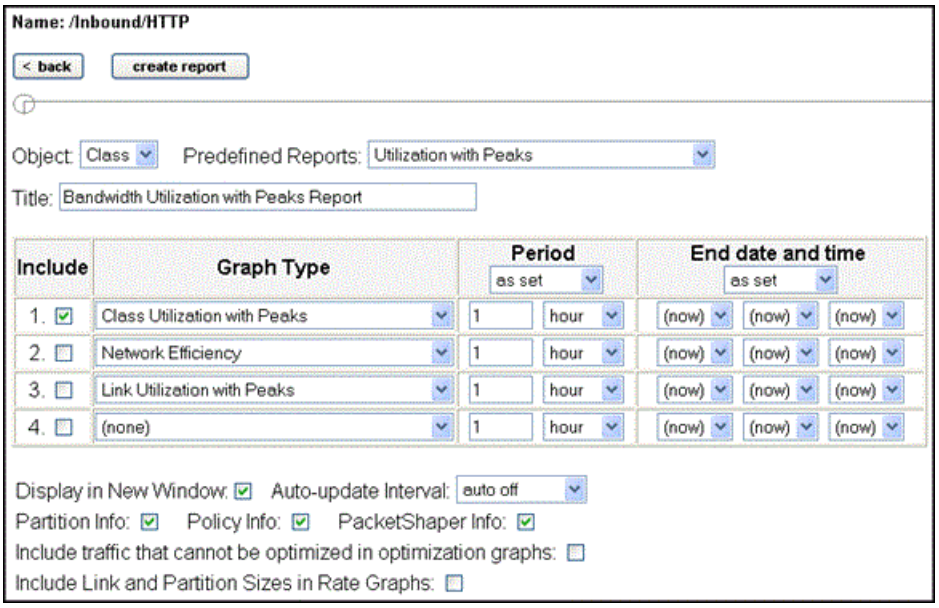

4. In the **Object** drop-down menu, select **Link**, **Partition**, or **Class**.

5. Select the graph name from the drop-down list in the Graph Type column. For descriptions and examples of all the graphs, see Preconfigured Graphs.

6. Specify the period you want to graph. Enter a number in the **Period** column and select the units (**hour**, **day**, **week**, or **month**).

7. Select an end date for the graph. The default end date is **now** (the current date). To change the date, select the month and day from the drop-down lists. You can change the date to any time within the past two months.

8. Select an end time for each graph. The default time is **now** (the current time). To change the time, select the time from the drop-down list.

9. Click **create report** to view the graph.

Analyzing Graphs

For information on analyzing graphs, see the following PacketGuide Recommendations:

Analyze an Application's Behavior
Identify Performance Saboteurs

Access Wasted Network Capacity
Analyze your Link's Behavior

Create Top 10 Categories Graph

Depending how your category classes are structured in your traffic tree, you may be able to create a graph of your top ten categories. If your category classes were auto-discovered or manually created under a single class (such as HTTP or a class with an "all categories" matching rule), you can create a Top 10 Children Classes graph of the parent class to see which categories are generating the most traffic on your network. For example, if you auto-discovered categories under an "all categories" class, you can select this parent class and graph its top ten children.

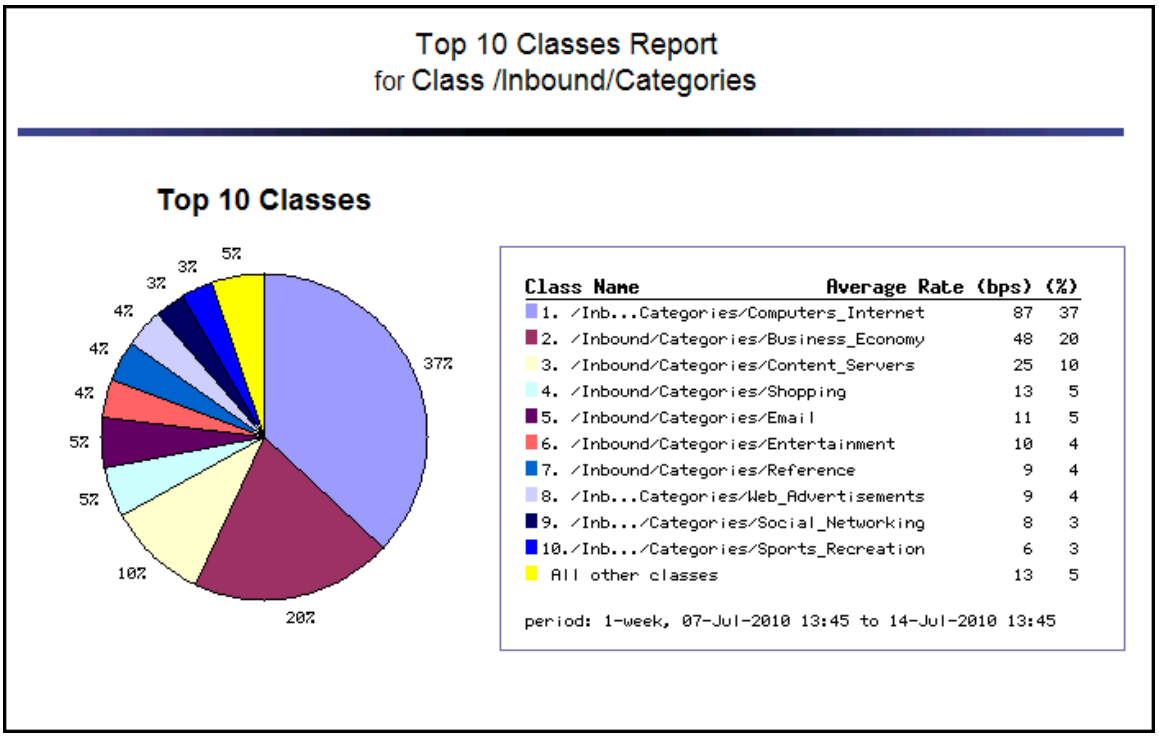

To create a top 10 categories graph:

1. Click the **Monitor** tab in the Advanced UI.

2. Click the icon next to the class that is the parent of the category classes. The *Statistics:Reports* screen appears.

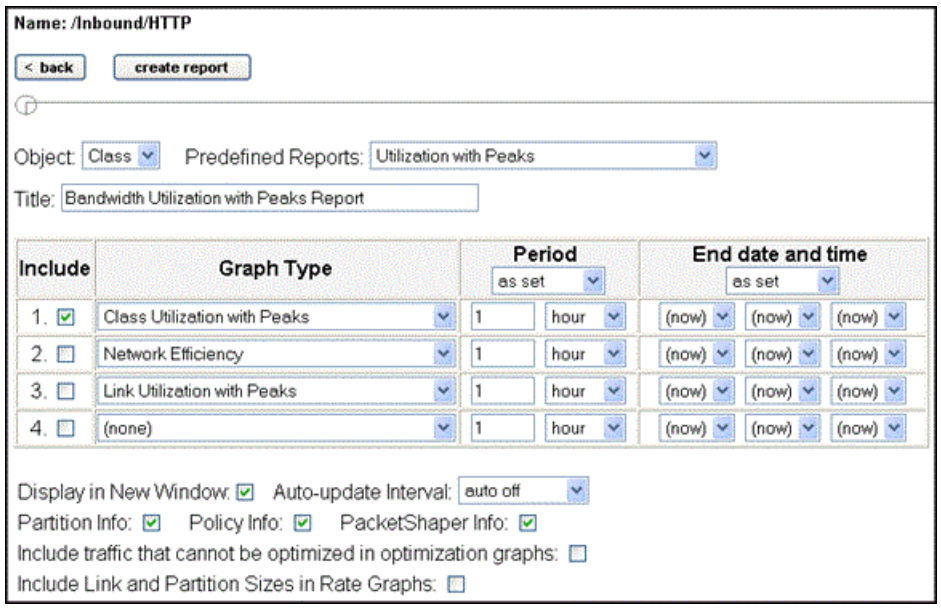

3. For **Graph Type**, select **Top 10 Children Classes**.

4. Click **create report** to display the graph with the default time period (last hour).

   or

   Specify the period you want to graph:

   - Enter a number in the **Period** column and select the units (**hour**, **day**, **week**, or **month**).
   - Select an end date for the graph. The default end date is **now** (the current date).
   - Select an end time for each graph. The default time is **now** (the current time).
   - Click **create report** to view the graph.

Graph a Specific Link, Partition, or Class

A report can include up to four graphs of a specific link, partition, or class. By selecting specific criteria, you can generate reports that answer questions such as:

- What was the most bandwidth the FTP traffic class consumed last week?
- Is a partition size appropriate for its traffic class?
- How efficient was my Oracle network traffic? What was its percentage of "good service?"
- What was responsible for my slow Citrix response yesterday a congested network or a sluggish server?

To graph a specific link, partition, or class:

1. In the **Manage** tab:

   - In the left window pane, select the link, partition, or class you want to graph.
   - Click **statistics** and choose **graph**. The *Statistics:Reports* screen appears. **show screen...**

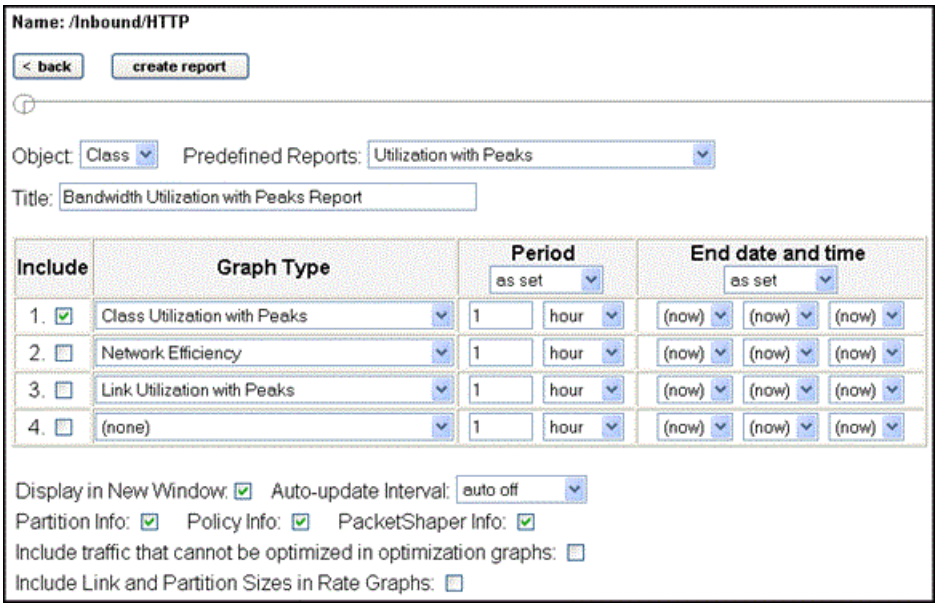

or

In the **Monitor** tab:

- Click the icon next to the class you want to graph. The *Statistics:Reports* screen appears.

2. In the **Object** drop-down menu, select **Link**, **Partition**, or **Class**.

   The object you select determines the report types available.

3. Select a report from the **Predefined Reports** drop-down list. The report types include:

- Utilization (the default)
- Utilization with Peaks
- Shaping Policy Effectiveness
- TCP Connections Initiated
- Connection Efficiency
- Top 10 (Partitions or Classes, depending on selected object)
- Custom

When you select any report type except **Custom**, the PacketShaper software enters graph suggestions in the center portion of the window; these suggestions are called the report's *member graphs*.

4. To label the report, enter a report title in the **Title** field.

5. To include a graph specified in the Graph Type column, select the **Include** checkbox. Uncheck the **Include** checkbox if you do not want to include the member graph.

6. To choose additional or different graph types, select the graph name from the drop-down list in the Graph Type column. For descriptions and examples of all the graphs, see Preconfigured Graphs.

> **Note:** You can select the same graph more than once to display information for different time intervals.

7. Determine whether you want to set the same or unique time intervals for the graphs in the report:

- With the default **Period** setting (**as set**), you can select a different period for each graph.
- To have each graph use the same time interval, select one of the "same as" options in the **Period** field. For example, choosing **Same as 1** uses the first graph's interval for the other graphs.

8. Enter a number in the **Period** column and select the units (**hour**, **day**, **week**, or **month**). The default is **1 hour** .

If you selected **as set** in step 9, set intervals for each included graph. If you selected one of the "Same as" options, you just need to set the interval for the graph you indicated.

9. Determine whether you want the same or unique **end date and time** settings for the graphs:

- With the default **End date and time** setting (**as set**), you can select a different end date and time for each graph.
- To have each graph use the same end date and time, select one of the "same as" options (**Same as 1**, **Same as 2**, and so on).

10. Select an end date for each graph. The default end date is **now** (the current date).

To change the date, select the month and day from the drop-down lists. You can change the date to any time within the past two months.

11. Select an end time for each graph. The default time is **now** (the current time).

To change the time, select the time from the drop-down list.

12. Click the **Display in New Window** checkbox to display the report in its own window. When this option is not checked, the report will display in the main browser window.

13. Determine how often you want the report data refreshed.

- Select an interval from the **Auto-update Interval** field to periodically refresh the data.
- Select **auto off** if you do not want the data automatically refreshed.

14. Determine the specification details you want the report to include by selecting or unchecking the respective checkboxes for **Partition Info**, **Policy Info**, and **PacketShaper Info**.

- **Partition Info** includes the partition minimum and maximum size.
- **PacketShaper Info** lists the IP address, name, and domain.
- **Policy Info** includes the policy type and settings.

> **Tip:** If you are generating and printing the same report from multiple units, be sure to select the **PacketShaperInfo** checkbox to differentiate between the reports.

15. If desired, select the **Include Link and Partition Sizes in Rate Graphs** checkbox. When this box is checked, the link or partition size will appear as a horizontal line at the top of link and partition utilization graphs.

16. Click **create report** to generate the report.

The report appears with included graphs. Recent traffic that was not stored in the latest data-sampling interval is not reflected in the graphs. **show screen...**

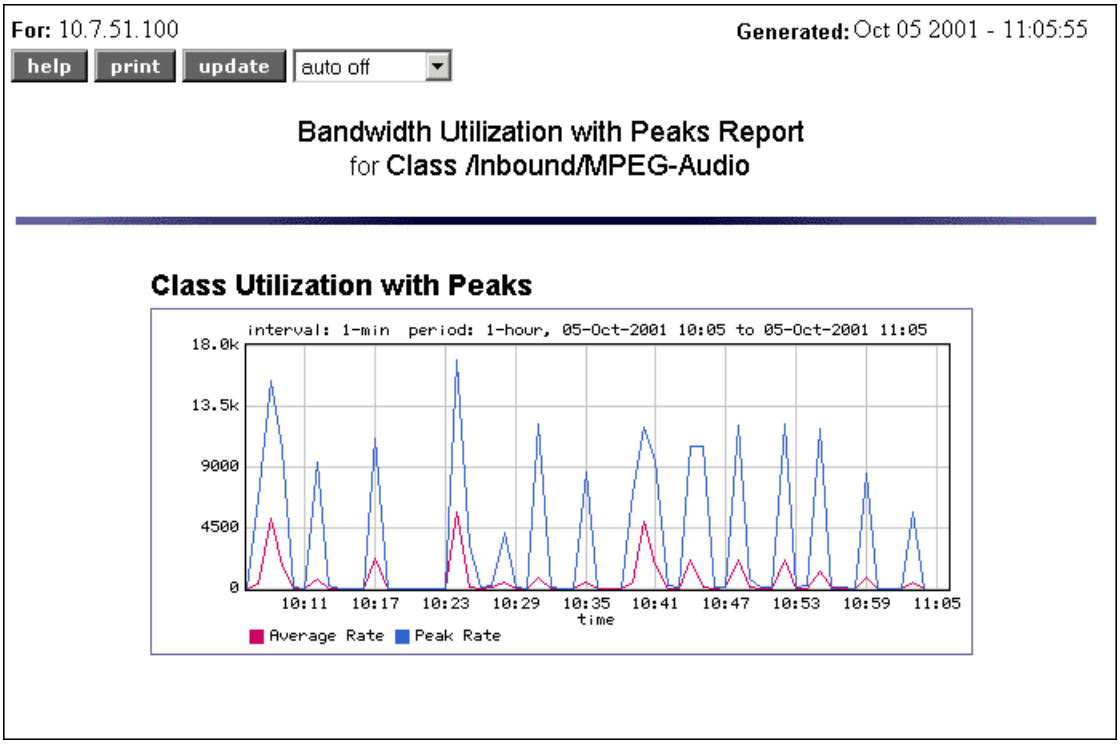

17. To print the report, click **print** at the top of the report window.

> **Note:** If a graph shows the partition size to be less than the actual size that was set, the unit is oversubscribed (that is, the sum of all partitions is greater than the link size). In the case of oversubscription, PacketShaper adjusts the programmed value to a scaled value, and uses the scaled value in the graph.

Additional Information

The traffic class you select for a report determines the graphical information you can see. Factors that influence the available reports and graphs are whether the class is a parent or a leaf class (a class with no children) and whether the class has an associated partition.

The object type (class, partition, or link) also affects the type of graphs available to you. For example, a partition offers a report and graph called "Top 10 Partitions," but a link does not. As another example, all three object types offer a Utilization report, but the reports contain different member graphs.

For information on analyzing graphs, see the following PacketGuide Recommendations:

Analyze an Application's Behavior
Identify Performance Saboteurs
Access Wasted Network Capacity
Analyze your Link's Behavior

Show Link Utilization

The Network Performance Summary is automatically displayed when you click the **Report** tab. This report offers a quick way to analyze link utilization. It displays graphs of the managed link's inbound and outbound traffic. **show screen...**

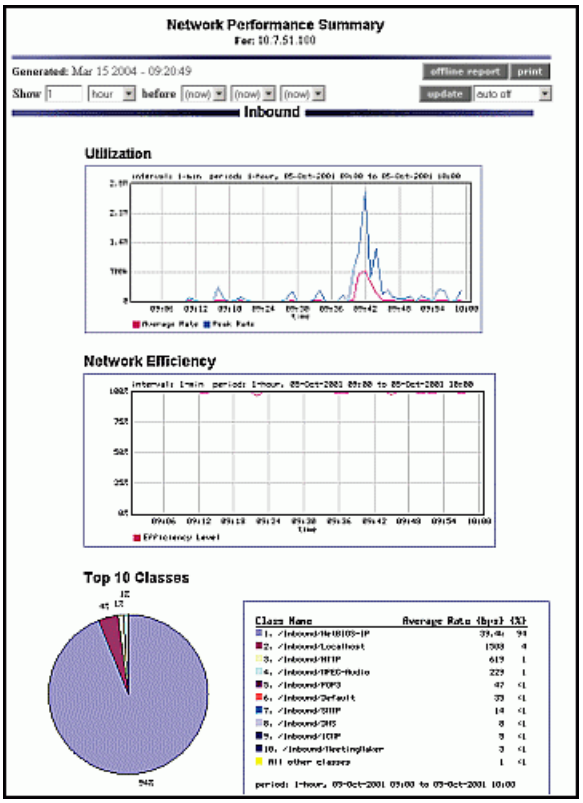

Specifically, the Network Performance Summary contains the following items:

- **Bandwidth Utilization** line graphs Shows the links average and peak bandwidth usage in bits per second
- **Network Efficiency** line graphs Shows the percentage of all bytes that are *not* TCP retransmits, giving you an indication of overall network health. For example, a value of 95% indicates that of all traffic on the network (TCP and UDP), 5% was retransmitted TCP traffic. (UDP is assumed to be 100% efficient since retransmissions are not calculated on this type of traffic.)
- **Top 10 Classes** pie charts Shows the relative portions of bandwidth allocated to the ten most active classes on the network; displays each class' average bandwidth usage in bits per second and its percentage of the total bandwidth on the link
- Total kilobytes sent (outbound) and received (inbound) during the interval

A blue horizontal line above the graphs indicate that traffic shaping was enabled during the graphing interval. This line shows you exactly when shaping was turned on so that you can easily see its effect on network traffic. For example, after enabling shaping, you can view and interpret the consequences of your policies and partitions. **show screen...**

Utilization

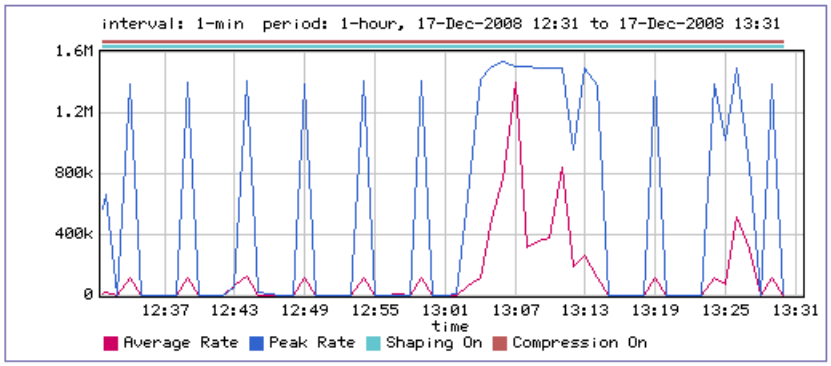

For information on analyzing this report, see the following PacketGuide Recommendations:

Analyze your Link's Behavior
Identify Performance Saboteurs
Access Wasted Network Capacity

The default report contains data from the current PacketShaper ("self"). If you have enabled Standby, you can view the Network Performance Summary of the other partners in the high availability cluster. Choose the partner's IP address from the **Standby Partner** list. For more information about the Standby feature, see "Configure Standby" on page 722.

By default, the Network Performance Summary graphs analyze network traffic over the last hour, but the time period and end date are user-customizable. See Selecting a Time Period.

Graph Multiple Classes

In the **Reports** tab, you can create reports that graph rate of multiple classes (aggregate length of class names must be less than 900 characters, 16 classes maximum). These reports can be saved and then retrieved at a later date. You can also have PacketShaper email the report to one or more recipients on a regular basis (for example, every day).

To design and save a report:

1. Click the **Report** tab.

2. Click **create report**. The *Create Reports* window appears. **show screen...**

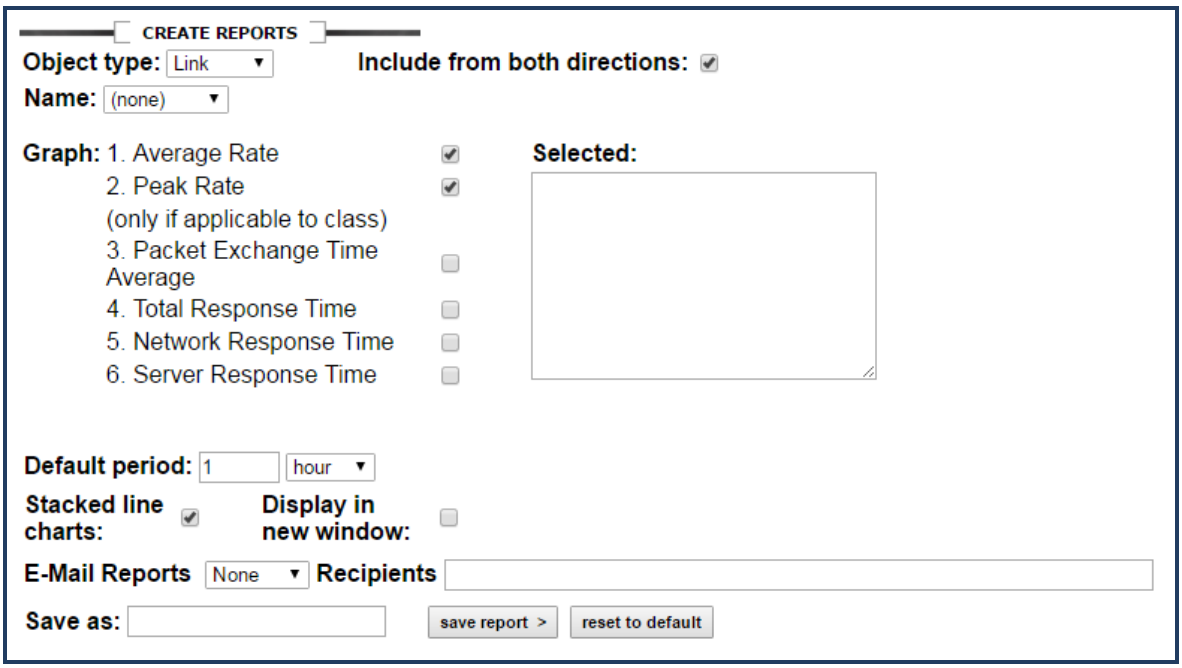

3. In the **Object type** field, select the type of object to be graphed (**Link**, **Partition**, or **Class**).

4. If you *do not* want to create the report for both inbound and outbound traffic, uncheck the **Include from both directions** check box.

5. Select the classes to be graphed (up to 16 classes):

   - For a Link object, select **(none)**, **/Inbound**, or **/Outbound**.
   - For a Partition or Class object, select a traffic class name.

6. Select the statistical information you want to graph, such as **Average Rate**.

> **Note:** Response-time statistics are available only for class objects for which RTM is available.

7. In the **Default period** field, enter a number and select **hour**, **day**, **week**, or **month**.

8. To produce separate line graphs for each of the selected objects, uncheck the **Stacked line charts** check box.

9. To display the report in a new browser window, select the **Display in new window** check box.

10. To have PacketShaper email the report on a regular basis:

    a. In the **Email Reports** field, select the frequency: **Hourly**, **Daily**, **Weekly**, **Monthly**.

    b. In the **Recipients** field, enter one or more email addresses of the individuals you want to receive the report. Although there is no limit to the number of recipients you can add, fewer than 10 is recommended.

> **Caution:** This feature requires that an email server be configured on PacketShaper. See " Configure an SMTP Server " on page 329.

11. To save the report for later retrieval, enter a name in the **Save as** field.

12. Click **save report**.

The name of your report will appear on a button in the left pane of the *report* window, under My Reports. **show screen...**

Perform Host Analysis

The Host Analysis reporting tool allows you to gain insight into host, user, and flow activity on your network. With this tool, you can display information formerly available only in the command-line interface using the hostdb info and traffic flow commands. Specifically, you can:

- List bandwidth utilization and flow information for IPv4 and IPv6 hosts or user names on the network. This list can consist of user names, inside hosts, outside hosts, active hosts, unknown hosts, or all hosts, and can be sorted by bandwidth utilization, new connections, failed connections, or host IP address/user name.
- Find flows for a particular IPv4 or IPv6 address, traffic class, service name, or port number. For each flow, the list provides the IP addresses of the conversation pair, the class name into which the PacketShaper classified the flow, and the protocol of the flow.
- Drill-down to find out detailed flow information for a suspicious host (for example, one that is using excessive bandwidth or creating an inordinate number of connections or failed connections).

The tool offers two tables that display a host's activity:

- Host Analysis Table
- Flow Analysis Table

Open the Host Analysis Tool

1. Click the **Report** tab.

2. In the Tools section on the left, click the **Host Analysis** button. The tool opens in a new window.

Host Analysis Table

The **Host Analysis** table displays information about hosts that have communicated through the PacketShaper.

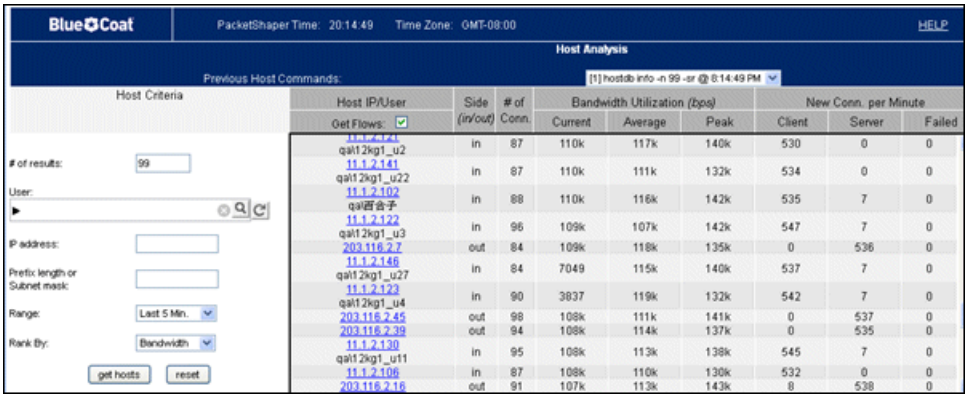

Enter information and make selections in the *Host Criteria* section to retrieve information about hosts into the *Host Analysis* table. If you don't specify any criteria, all hosts are displayed in the Host Analysis table. To narrow down the results, you can specify criteria, such as a user name, IP address, or subnet. If more than one criteria is specified, a host must match all criteria in order for it to be listed in the Host Analysis table.

1. Specify the **# of results** to display. Default is 99.

2. To analyze host information for a particular user, specify the **User** name using any of the following methods. The *Host Anlaysis table* will list all IP addresses associated with the user name. (Requires that BCAAA be configured. See Configure PacketShaper as a BCAAA Client.)

- Enter all or any part of the user name (such as the first or the last name) in the **User** field. For example, enter *john*, *smith*, or *john.smith*. Then click the Search icon and select the name from the search results list.
- To select the user name from a list all users, leave the **User** field empty and click the icon; then select the name from the search results list.
- To list all user names for a specific domain, enter the domain name and click the icon; then select the name from the search results list.
- Manually type the domain name and user name, for example, *CAL\john.smith*. There is no need to click the icon unless you want to verify that the user name is correct.

3. Specify the **IPaddress** (IPv4 or IPv6) of the target host to limit the results to a specific host. Default is blank (retrieves all hosts).

4. Specify the **Prefix length orSubnet mask** to limit results to a specific IPv4 or IPv6 subnet. Default is blank (retrieves all subnets). For example, *255.255.255.0* for an IPv4 subnet mask or *64* for an IPv6 prefix length.

5. Select a **Range** to limit the results to a particular type of activity:

- **Inside only**: displays hosts located on the inside of the PacketShaper
- **Outside only**: displays hosts located on the outside of the PacketShaper
- **Unknown**: displays only those hosts whose location relative to the PacketShaper is unknown
- **Active only**: displays currently active hosts
- **Last 5 min.**: displays all hosts with activity within the last five minutes (default)
- **All**: displays all hosts that have communicated through the PacketShaper

6. Choose to **Rank By** results by:

- **Bandwidth**: orders hosts by bandwidth used
- **New Conn.**: orders hosts by number of new connections
- **Failed Conn.**: orders hosts by number of failed connections
- **Host IP**: orders hosts by ascending IP address

7. Click **get hosts** to retrieve data into the **Host Analysis** table.

Additional Information:

- To restore the *Host Criteria* to their default settings, click **reset**.
- Each time you click **get hosts**, the results are saved on the PacketShaper. You can access the results from recent host queries by making a selection from the **Previous Host Commands** drop-down list, located above the Host Analysis table.
- To clear all data from the Host Analysis and Flow Analysis tables, refresh the browser window by clicking Refresh on the browser's toolbar or by pressing F5.

The following table describes the data displayed in the **Host Analysis** table.

Column	Description
Host IP/User	IPv4 or IPv6 addresses of hosts that have communicated through the PacketShaper; if the User Awareness feature is configured, this column also displays user names. When **Get Flows** is enabled (default), clicking a host's IP address populates the Flow Analysis table with information about traffic flows for the specified host or user name.
Side (*in/out*)	Indicates the location of the host (inside or outside) relative to the PacketShaper.
# of Conn.	The number of connections that were active when the **get hosts** button was clicked
Bandwidth Utilization (*bps*)	The current, average, and peak bandwidth utilization of the host or user, as measured in bits per second (bps). *Current bps*: the current rate for the host or user *Average bps*: the one-minute moving average for the host or user rate *Peak bps*: the highest rate the host's or user's connection has reached. This is the sum of the inbound and outbound traffic, relative to the host or user.
New Conn. per Minute	The number of new connections per minute. Depending on the application, a host can act as either a Client or a Server. *Client*: The number of new client application connections per minute made by this host *Server*: The number of new server application connections per minute made by this host *Failed*: The number of failed connections reported by the host

Flow Analysis Table

The **Flow Analysis** table displays traffic flow information for TCP connections and/or UDP sessions.

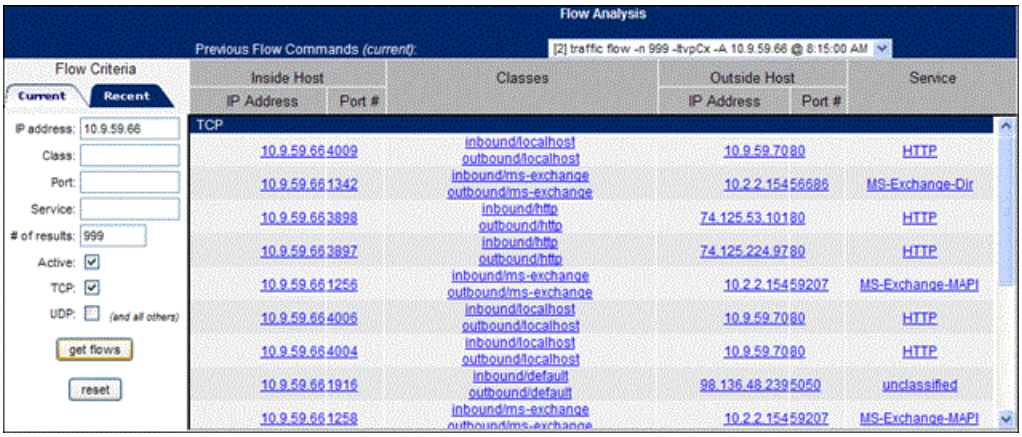

Retrieve Current Flows

Enter information and make selections in the **Current** tab on the *Flow Criteria* section to retrieve information about current flows (that is, flows that have been active within the past 60 seconds).

1. Specify the **IPaddress** (IPv4 or IPv6) of the target host to limit the results to a specific host. Default is blank (retrieves flows for all hosts).

2. Specify the **Class** name to limit results to a specific class, such as *Inbound/Default*. Default is blank (retrieves flows for all traffic classes).

3. Specify a **Port** number to limit results to flows on a specific port. For example, enter **80** to retrieve only HTTP flows.

4. Specify the exact **Service** name to limit results to flows of a specific service, such as **HTTP**.

5. Specify the **# of results** to display. Default is 999.

6. Select the **Active** checkbox to display only active flows. Default is checked.

7. Select the **TCP** checkbox to display only TCP flows. Default is checked.

8. Select the **UDP** checkbox to display UDP flows in addition to TCP flows. Default is unchecked (UDP flows not displayed).

9. Click **get flows** to retrieve data into the **Flow Analysis** table.

> **Tip:** To view flows for a specific host, click the host's IP address in the **Host Analysis** table.

Retrieve Recent Flows

In addition to viewing current flows, you can also view recent flows (that is, flows that occurred in the last hour but that are not currently active).

1. Click on the **Recent** tab on the *Flow Criteria* section of the **Flow Analysis** table.

2. Specify a class name, IP address, or DNS name as the **Target** for flow retrieval.

3. Click **get flows** to retrieve data into the **Flow Analysis** table.

Additional Information:

- Each time you click **get flows**, the results are saved on the PacketShaper. You can access results of recent flow queries by making a selection from the **Previous Flow Commands** drop-down list, located above the Flow Analysis table.
- To restore the Flow Criteria to their default settings, click **reset**.

The following table describes the data displayed in the **Flow Analysis** table.

Column	Description
Inside Host	Displays information about the inside host of the flow *IP Address*: the IPv4 or IPv6 address of the inside host. Click the IP address to view all flows on that inside host. *Port #*: the port number of the inside host used by the flow. Click the port number to view all flows on that port.
Classes	Displays the inbound and outbound traffic class for the traffic flow
Outside Host	Displays information about the outside host of the flow *IP Address*: the IPv4 or IPv6 address of the outside host. Click the IP address to view all flows on that outside host. *Port #*: the port number of the outside host used by the flow. Click the port number to view all flows on that port.
Service	The service associated with the traffic flow. Click the service name to view all flows for that service.

Create Statistical Reports

The Statistical Graphing Tool provides a simplified method for graphing the measurement data that is associated with each traffic class. Because the measurement data is preloaded into the tool whenever you select a class from the Class Tree, graphs based on the selected class are rendered immediately. You can even scroll through your traffic tree, viewing graphs as you go. This allows you to explore the traffic statistics for your applications quickly and easily and reveal relationships that you may not have noticed before.

The Statistical Graphing Tool provides three views that enable you to display graphical information about a particular class:

- **Historical Overview**—Allows you to graph up to four measurement variables for a selected class over the last hour, day, or week.
- **Reports**—Allows you to generate reports that show the values of a set of measurement variables graphically. You can create standard reports, which contain preconfigured graphs. Or, you can create custom reports, which graph a set of measurement variables that you define.
- **Real-Time Graphing**—Displays the current bandwidth utilization (in bits per second) for the selected class.

To access the Statistical Graphing Tool:

1. Click the **Report** tab.
2. In the Tools section on the left, click **Statistical Graphing**.

> **Note:** In order to use the Statistical Graphing Tool, you must have the Adobe Flash Player plug-in. If you do not have it installed, you will be prompted to download and install it before you can proceed.

3. In the Class Tree, select the class you want to graph. As soon as you click on a class name, the corresponding graph is displayed in the right window pane. **show screen...**

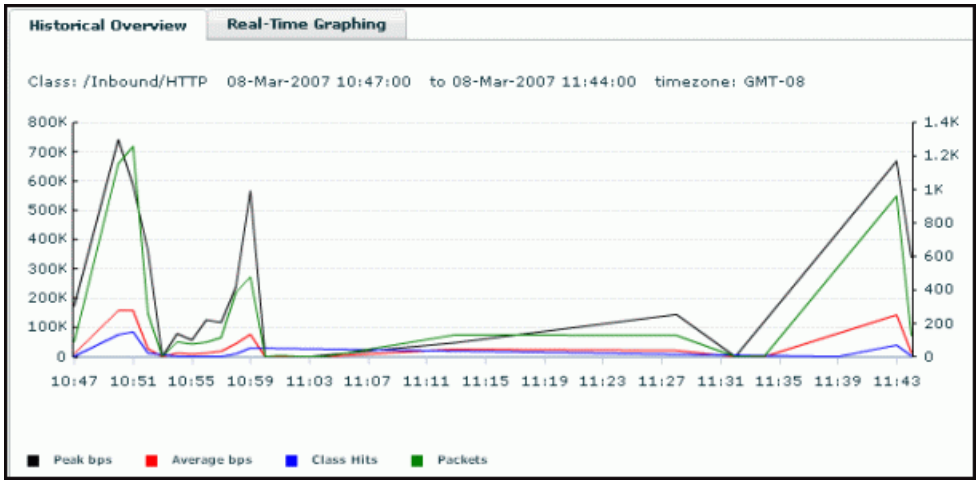

Click the ▶ arrow next to the Inbound or Outbound folder to expand the tree. If you have many classes to navigate through,

you can filter the class tree by entering a search string in the search field. Only classes containing the string will be displayed. You can even use the up and down arrow keys on your keyboard to scroll through the classes, viewing graphs as you go.

4. Define the measurement variables you want to graph by selecting variable descriptions from the **Line 1**, **Line 2**, **Line A**, and **Line B** drop-down lists. The set of measurement variables available in this view enable you to quickly assess if there is a problem with the class.

By default, the Total Bytes variable (Line 1) is graphed. You can select up to four variables to graph on the Historical Overview. Note that the variables that you select for Line 1 and Line 2 will be graphed on the Y-axis on the left side of the graph and the variables you select for Line A and Line B will be graphed on a second Y-axis on the right side of the graph. This allows you to graph variables that are measured on different scales on the same graph. For example, suppose you wanted to graph Total Bytes, Packets, Network Delay, and Packet Round Trip. Total Bytes and Packets are both measured as an integer and it therefore makes sense to graph them on the same axis. Likewise, Network Delay and Packet Round Trip are both measured in milliseconds, so it makes sense to graph them on the other axis.

5. Define the time period for the graph.

By default, data from the last hour is graphed. You can set the time period by selecting **Week**, **Day**, or **Hour**. When you change the time period, the new graph renders instantly.

6. To see an actual value associated with a specific point in time on the graph, mouse over the point. For example, mouse over a point on the Total Bytes line to see a pop-up tool tip, such as Total Bytes 16:12 169347.

7. If you leave the Historical Overview graph open for a while and then come back to it later, click **Update** to generate an up-to-date graph.

8. To print the graph that is currently displayed in the Historical Overview window, click **Print**.

> **Note:** Do not open more than one instance of the Statistical Graphing Tool.

Generate VoIP Reports

You can view performance metrics for VoIP calls on the VoIP Summary report.

To view the VoIP Summary report:

1. In the Advanced UI, click **VoIP Summary**.
2. In the **Class** drop-down, select the VoIP class you want to report on and click **update**.
3. By default, reports show metrics for the previous hour. To view metrics for other intervals and time periods, do the following:
    - For the previous eight hours, day, or week, click the corresponding button (for example, **Last 8 hr**).
    - For a specific time period, choose a duration from the **Show** drop-down lists and an end point from the **before** drop-down lists, and then click **update**.
    - To automatically update the VoIP Report page, select an interval (for example, **every 5 min**) from the update drop-down list, and then click **update**.
4. The default report contains data from the current PacketShaper. If you have enabled Standby, you can view the VoIP Summary of the other partners in the high availability cluster. Choose the partner's IP address from the **Standby Partner** list. For more information about the Standby feature, see "Configure Standby" on page 722.

Up to seven reports are available:

- **VoIP Latency***:* the time required for packets to travel from one PacketShaper to another, in milliseconds (note that latency is calculated on the Inbound interface only). Latency calculations are supported by synthetic RTP packets (probe requests) which are issued by PacketShapers at a configurable interval when an RTP flow is active. The header of the synthetic RTP packet contains a timestamp, sequence number, and unit ID information which are used to measure the round-trip time between the two PacketShapers. For instance, if the sum of measured latency is 500 ms and the number of latency samples is 25, then the average latency is 20 ms.
- **VoIP Jitter***:* the variation in the delay of received packets in a flow. Jitter is measured by comparing the interval when RTP packets were sent to the interval at which they were received. For instance, if packet #1 and packet #2 leave 50 milliseconds apart and arrive 60 milliseconds apart, then the jitter is 10 milliseconds.
- **VoIP Loss**: the percentage of lost packets. Calculation of packet loss (that is, expected packets that are not received) is performed per RFC 3550 using RTP header sequence numbers.
- **MOS multiplied by 10** : the Mean Opinion Score, which provides a relative measurement of VoIP call quality. PacketShaper measures MOS using a scale of 10 (bad) to 50 (excellent); scores of 35 or lower indicate possible problems with call quality. To convert to a standard MOS score (which uses a scale of 1-5), divide the PacketShaper MOS value by 10.
- **R-Factor**: the transmission quality rating, with a typical range of 50-100. An R-Factor of 80 or higher is highly desirable; scores of 70 or lower indicate possible problems with call quality.
- **Class Hits** and **Peak Flows**: When an RTP-I class is specified in the **Class** field, two additional reports are displayed: Class Hits and Peak Flows. These reports display the numbers of class hits and concurrent flows for the specified time period.

Additional Information

- To enable latency, R-Factor, and MOS measurements of VoIP flows between PacketShapers, you must first change each unit's system variable **Enable Latency Calculations** from **off** (default) to **on**. See Enabling Latency Calculations.
- Because VoIP reports are derived from VoIP metrics, they are subject to the same requirements and limitations. See VoIP Metrics Requirements.
- The amount of historical data available depends on the capacity of your PacketShaper and the volume of network traffic.
- The Class Hits and Peak Flows reports are only displayed when an RTP-I class is specified in the **Class** field.

VoIP Latency Calculation

When latency calculations are enabled on a PacketShaper, an active RTP flow triggers the PacketShaper to issue synthetic RTP packets at a configurable interval (default is five seconds). The headers of synthetic RTP packets (also known as probe requests) contain a timestamp, sequence number, and unit ID information, which are used to measure the round-trip time between PacketShapers.

The following illustration shows the packet flow of a VoIP call when latency calculations are enabled. Note that latency is measured between the two PacketShapers, not between the two VoIP phones.

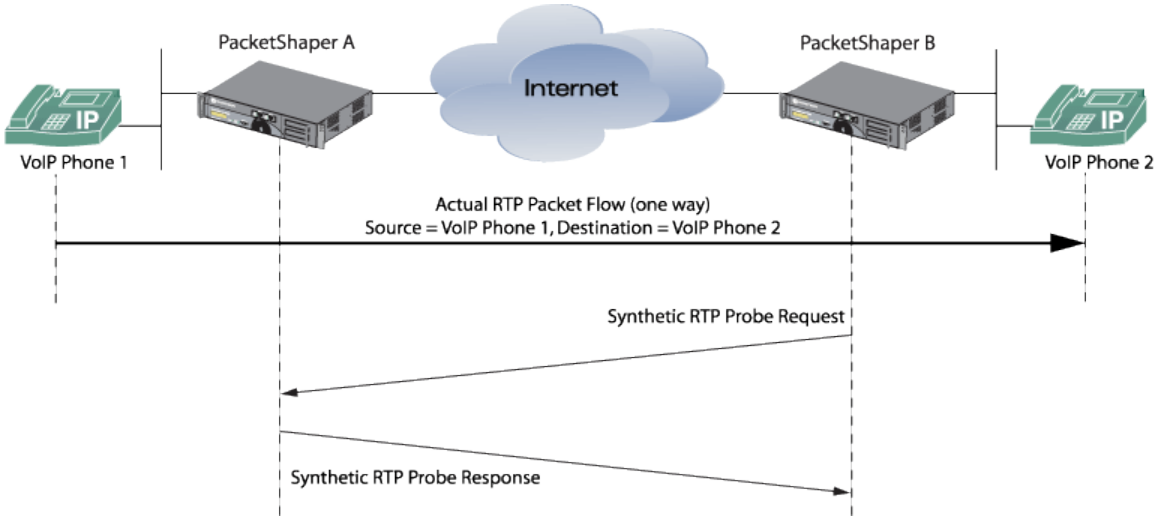

VoIP R-Factor and MOS

Rating Factor (R-Factor) and Mean Opinion Score (MOS) are two commonly-used measurements of overall VoIP call quality.

- **R-Factor**: A value derived from metrics such as latency, jitter, and packet loss per *ITU-T Recommendation G.107*, the R-Factor value helps you quickly assess the quality-of-experience for VoIP calls on your network. Typical scores range from 50 (bad) to 90 (excellent).
- **MOS**: A value are derived from the R-Factor per *ITU-T Recommendation G.10* which measures VoIP call quality. PacketShaper measures MOS using a scale of 10-50. To convert to a standard MOS score (which uses a scale of 1-5), divide the PacketShaper MOS value by 10.

Correlations between R-Factor and MOS scores and the user's experience of VoIP call quality are shown in the following table (standard MOS values are shown in parentheses):

R-Factor	MOS	User Experience
90	43 (4.3)	Excellent
80	40 (4.0)	Good
70	36 (3.6)	Fair
60	31 (3.1)	Poor
50	26 (2.6)	Bad

Measurement Statistics Tasks

PacketShaper stores data on dozens of measurement variables, such as avg-round-trip-time and class-hits. For classes and partitions, units store at least one day's worth of per-minute data samples and at least one month's worth of per-hour data samples. For links, units store seven days of per-minute data samples and six months of per-hour samples.

This section contains tasks for viewing measurement data on the screen, downloading the data to a file, and tracking host data.

View Measurement Statistics

PacketShaper's statistical reports are user customizable. You specify the object for which you want to gather measurement data (a specific link, partition, or class), and choose which type of data you want to analyze such as class hits, number of TCP connections refused by the server, packet count, or count of retransmissions. There are dozens of measurement variables from which you can choose. (See Measurement Variables for a complete list.) You also specify the time period for which you want to view statistics (such as for the last week or previous two months) and the increments for the data output (such as 15-minute increments).

The measurement data output is a comma-delimited list. You can display this list in your browser window to preview it before importing it into a spreadsheet or database application. Alternatively, you can export the measurement data to a file.

To display a list of measurement statistics in comma-delimited format:

1. Click the **Manage** tab.

2. Select the link, partition, or class for which you want to view statistics.

3. Click **statistics** and choose **data**. The *Statistics:Data* window appears. **show screen...**

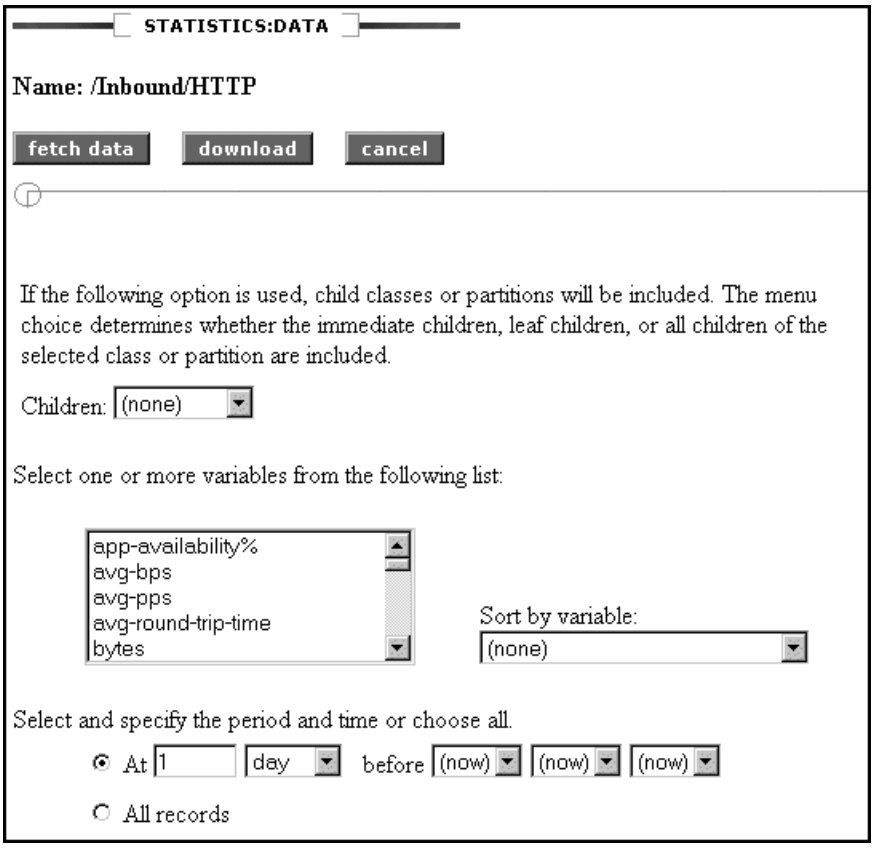

4. If you are asked for the type of object, select **link**, **partition**, or **class**.

5. If you selected partition or class, you also can select one of the **Children** options: (If you don't see a Children field, the class you selected has no child classes.)

none data for child classes is *not* included in the output
immediate includes data for the specified class and direct child classes (not grandchildren classes)
leaf includes data for leaf classes only (that is, child classes that don't have any children)
all includes data for the complete branch: the specified class, direct child classes, and all other classes that are descendants of the child classes.

PacketShaper lists only the measurement variables that are appropriate for the object you selected. Some variables are common to all object types, while others are specific to links, partitions, or classes.

6. Select the variables to be viewed:

- To select multiple variables, hold down Ctrl as you click each variable name.
- For variable descriptions, see Measurement Variables.

> **Note:** Some values are recorded for either a parent or a child class, but not for both. If you select a variable that has no recorded measurements for the selected object, a zero will be displayed in the output.

7. Select the variable to sort by:

- Be sure that the sort variable is one that you selected in step 3.

  or

- Choose **none** to dump variables in alphabetical order.

8. Set the time period:

- Choose **All Records** to display all stored data. For classes and partitions, units store at least one day's worth of per-minute data samples and at least one month's worth of per-hour data samples. For links, units store seven days of per-minute data samples and six months of per-hour samples.

  or

- In the **At** field, enter the number of minutes, hours, days, weeks, or months for which you want data, and select the unit of time. Then specify the date and time of the end of the collection period. For example, choosing a *1 week* period with an end date and time of *Jan 1 0:00* indicates that you want a data dump covering the last week of December. Note that the "now" setting for date and time dumps the most recent data.

9. Select a data format. Aggregate formats show totals for the selected variables.

- **Time series** listed with the most recent time first
- **Time series - ordered by objects** listed in link, partition, or class order
- **Aggregate** summarized or totaled
- **Aggregate - ordered by objects** summarized in link, partition, or class order
- **Aggregate - ordered by variables** summarized in variable name order
- **XML for Oracle 8i** returns data in XML format that can be imported directly into Oracle 8i

> **Note:** The meaning of "total" depends on the variable. In some cases, "total" is not meaningful. For example, the total for the *bytes* variable is truly a total an accumulation of the minute counts for the *bytes* variable; however, for variables such as *avg-bps* and *tcp-retx-pkts%*, the totals are derived values.

10. Select a data sample interval. If you choose an interval less than one hour, minute samples are used. If you choose an interval greater than one hour, hourly samples are used. For example, if you requested one hour of data and you want it reported in 15-minute increments, choose **15 minutes**. PacketShaper will sum its samples, if necessary, to report the correct totals for your interval.

11. Click **fetch data** to display the comma-separated values in your browser window. **show screen...**

```
"link:/Inbound"
"time","avg-bps","bytes"
"03-Oct-2001 13:12:00",66704,500282
"03-Oct-2001 13:11:00",27607,207050
"03-Oct-2001 13:10:00",41959,314696
"03-Oct-2001 13:09:00",2323,17426
"03-Oct-2001 13:08:00",98,735
"03-Oct-2001 13:07:00",381,2859
"03-Oct-2001 13:06:00",85,641
"03-Oct-2001 13:05:00",123,924
"03-Oct-2001 13:04:00",261,1960
"03-Oct-2001 13:03:00",98,735
"03-Oct-2001 13:02:00",104,782
"03-Oct-2001 13:01:00",216,1617
"03-Oct-2001 13:00:00",98,735
"03-Oct-2001 12:59:00",104,782
"03-Oct-2001 12:58:00",241,1805
"03-Oct-2001 12:57:00",85,641
"03-Oct-2001 12:56:00",85,641
"03-Oct-2001 12:55:00",275,2063
"03-Oct-2001 12:54:00",112,843
"03-Oct-2001 12:53:00",148,1112
"03-Oct-2001 12:52:00",381,2859
"03-Oct-2001 12:51:00",92,688
"03-Oct-2001 12:50:00",85,641
"03-Oct-2001 12:49:00",334,2503
"03-Oct-2001 12:48:00",104,782
"03-Oct-2001 12:47:00",1032,7743
"03-Oct-2001 12:46:00",172,1288
"03-Oct-2001 12:45:00",3169,23764
"03-Oct-2001 12:44:00",862,6466
"03-Oct-2001 12:43:00",894,6707
"03-Oct-2001 12:42:00",3689,27667
```

12. The fetched data appears in a new browser window. To save the data in a text file, use the browser's **File > Save As** command.

The amount of measurement data depends on the number of classes in your traffic tree. Typically, if the unit has been running uninterrupted, at least one day of data is available for intervals less than one hour and at least one month of data is available for intervals greater than one hour.

Download Measurement Statistics to a File

You can save measurement data in a comma-separated-value (.csv) file that can be opened with a spreadsheet application, such as Microsoft Excel, or imported into a database application.

To save measurement data to a file:

1. Click the **Manage** tab.

2. Select the link, partition, or class for which you want to download statistics.

3. Click **statistics** and choose **data**. The *Statistics:Data* window appears. **show screen...**

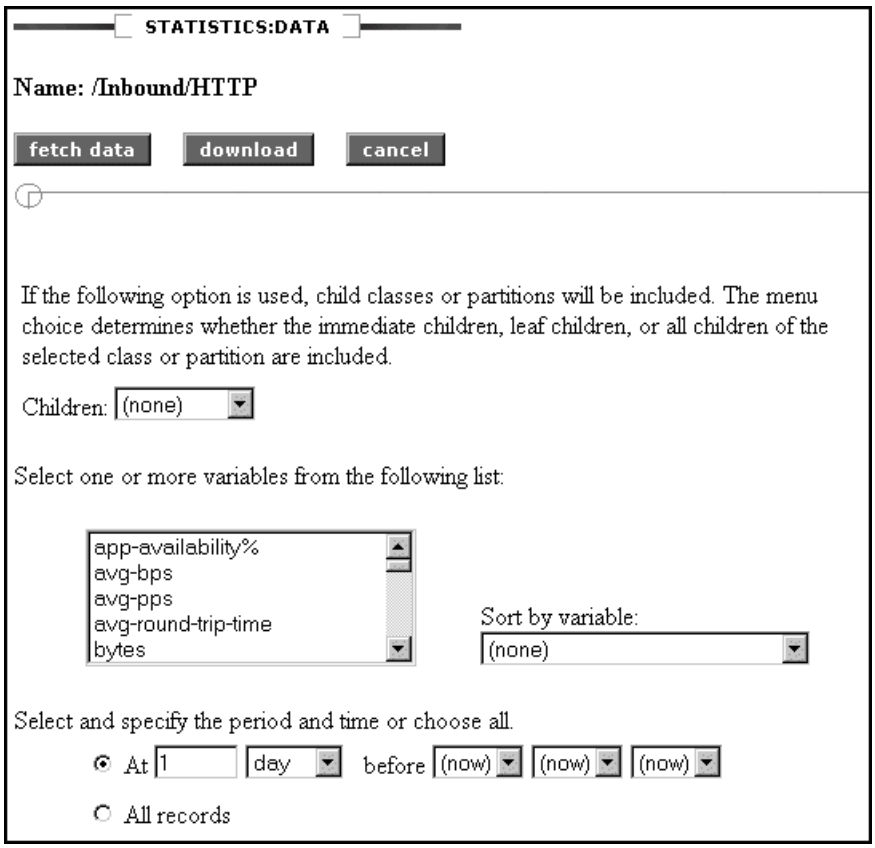

4. If you are asked for the type of object, select **link**, **partition**, or **class**.

5. If you selected partition or class, you also can select one of the **Children** options: (If you don't see a Children field, the class you selected has no child classes.)

none data for child classes is *not* included in the output
immediate includes data for the specified class and direct child classes (not grandchildren classes)
leaf includes data for leaf classes only (that is, child classes that don't have any children)
all includes data for the complete branch: the specified class, direct child classes, and all other classes that are descendants of the child classes.

PacketShaper lists only the measurement variables that are appropriate for the object you selected. Some variables are

common to all object types, while others are specific to links, partitions, or classes.

6. Select the variables to be exported:

- To select multiple variables, hold down Ctrl as you click each variable name.
- For variable descriptions, see Measurement Variables.

> **Note:** Some values are recorded for either a parent or a child class, but not for both. If you select a variable that has no recorded measurements for the selected object, a zero will be displayed in the output.

7. Select the variable to sort by:

- Be sure that the sort variable is one that you selected in step 3.

  or

- Choose **none** to dump variables in alphabetical order.

8. Set the time period:

- Choose **All Records** to export all stored data. For classes and partitions, units store at least one day's worth of per-minute data samples and at least one month's worth of per-hour data samples. For links, units store seven days of per-minute data samples and six months of per-hour samples.

  or

- In the **At** field, enter the number of minutes, hours, days, weeks, or months for which you want data, and select the unit of time. Then specify the date and time of the end of the collection period. For example, choosing a *1 week* period with an end date and time of *Jan 1 0:00* indicates that you want a data dump covering the last week of December. Note that the "now" setting for date and time dumps the most recent data.

9. Select a data format. Aggregate formats show totals for the selected variables.

- **Time series** listed with the most recent time first
- **Time series - ordered by objects** listed in link, partition, or class order
- **Aggregate** summarized or totaled
- **Aggregate - ordered by objects** summarized in link, partition, or class order
- **Aggregate - ordered by variables** summarized in variable name order
- **XML for Oracle 8i** returns data in XML format that can be imported directly into Oracle 8i

> **Note:** The meaning of "total" depends on the variable. In some cases, "total" is not meaningful. For example, the total for the *bytes* variable is truly a total an accumulation of the minute counts for the *bytes* variable; however, for variables such as *avg-bps* and *tcp-retx-pkts%*, the totals are derived values.

10. Select a data sample interval. For example, if you requested one hour of data and you want it reported in 15-minute increments, choose **15 minutes**. PacketShaper will sum its samples, if necessary, to report the correct totals for your interval.

11. Click **download**. A window opens that prompts you to save the data to a file.

12. Click **save file**.

13. Specify the path and filename, including .csv extension, in the browser's Save window.

> **Note:** The browser default may be set to save the file as "cgi.html," so you must explicitly set the filename extension to .csv.

PacketShaper records the statistics' dates and times in a format that may not match your application's standard (for example, *2/2/1999 4:47:19 AM*). The dates can be converted within the application. In Excel, you can increase the width of the date column to accommodate this format.

To graph the data using Microsoft Excel:

1. Open the .csv data file in Excel. The Excel application will detect that the file is a text file that must be converted from comma-delimited data into Excel format.

2. Use Excel's Chart Wizard to create a graph.

Back Up Measurement Data

If you back up your measurement data files, you can use the backups in case data later becomes corrupted. You may also want to create a backup to archive all the current data on a unit.

PacketShaper keeps track of measurement data according to the type of measurement variable (link, partition, or class) and type of sample (day or month). Day samples, measurement data that was recorded in one-minute intervals, are stored for at least a day. Month samples, data that was recorded every hour, are stored for at least a month. Thus, measurement data is stored in groups: link day, link month, partition day, partition month, class day, and class month. If you are using the host accounting feature, an additional measurement group is available for this type of data.

To back up all measurement data:

1. Access the CLI.
2. Stop the recording of measurement data: measure stop
3. Go to the Advanced UI in Firefox or Internet Explorer.

> **Note:** Some file operations do not work in the Chrome browser so it's best to use Firefox or IE.

4. Open the File Browser.
5. Go to the 9.258/me/measure folder and select the data files.
6. Zip the selected measurement data files.
7. Restart the recording of measurement data: measure start
8. Use the File Browser to download the zip file to a local workstation or network drive.

Restore Measurement Data

After you have backed up measurement data, you can restore it in case data became corrupted after the last backup or if you want to copy measurement data to another unit.

If you are restoring measurement data files to a different unit from which they were backed up, note the following:

- The unit to which you are restoring the data must be the same model, use the same software version, and have the same specifications as the unit from which you backed up the data.
- It is recommended that you back up all the measurement data on the target unit before copying over the new files.

To restore all measurement data:

1. Access the CLI.
2. Stop the recording of measurement data: measure stop
3. Go to the Advanced UI in Firefox or Internet Explorer.

> **Note:** Some file operations do not work in the Chrome browser so it's best to use Firefox or IE.

3. Open the File Browser.
4. Delete the files in the 9.258/me/measure folder.
5. Upload the backed up zip file of the measurement data to the 9.258/me/measure folder.
6. Unzip the file.
7. In the CLI, restart the recording of measurement data: measure start

Reset Measurement Data

When you reset measurement data, the data stored in the measurement database is cleared. It's necessary to reset measurement data after:

- upgrading PacketShaper software, in order to use any new measurement variables introduced in the new software
- loading a license that affects the number of classes

You have the option of resetting all measurement data or certain types of data: link, partition, or class .

To reset the measurement data:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, choose **unit resets**. The Unit Resets options appear on the *Setup* screen. **show screen...**

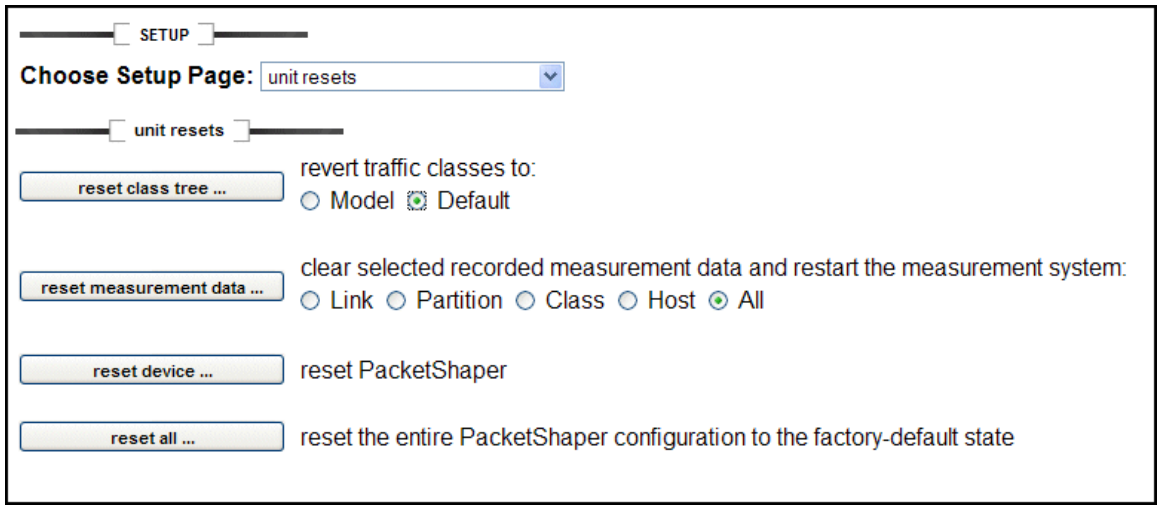

3. Select the type of measurement data to reset: **Link**, **Partition**, **Class** or **All**.
4. Click **reset measurement data**.

After you issue the command, you are prompted to confirm your reset request. Accumulated measurement data is cleared and the unit resets. Measurement and reporting data will not be available for several minutes.

Following the reset, do *not* attempt to load a new software image, as the file transfer will conflict with measurement engine processing. You can use the **measure show** command in the CLI to check the status of the measurement engine.

Host Accounting

The Statistics Data screen (see View Measurement Statistics) allows you to track bandwidth usage by traffic class. But if you do not have classes defined for each host (for example, when you are using dynamic partitions) and you need to track usage by host, you have another avenue available: host accounting. The host accounting feature offers a way to track byte counts of throughput for each IP address, group of users in a host list, or subnet. Here are several ways you could use this feature:

- Departmental billing (to determine how much bandwidth each department is using and then bill each department according to usage)
- Quota-based billing (to charge a certain rate for one usage level and then charge a different rate for higher usage)
- Service-level-based billing (to have different rates for different service levels for example, gold service pays more than silver for same usage)
- Group billing (to roll up individual usage figures into grouped totals)
- Cutoffs (to restrict access after a user exceeds prescribed bandwidth)

Follow these basic steps to use the host accounting feature:

1. Define the host accounting categories.
2. Enable host accounting.
3. Retrieve host accounting data.

Measure VoIP Statistics

The quality of the user experience for Voice over Internet Protocol (VoIP) applications depends greatly on the performance of the network connection between the calling parties. The VoIP metrics feature allows you to measure five metrics that reflect VoIP quality:

- *Jitter*–the variation in the delay of received packets in a flow
- *Latency*–the time required for packets to travel from one PacketShaper to another, as measured by the formula round-trip time / 2. Note that latency is calculated on the Inbound interface only.
- *Packet loss*–the percentage of lost packets
- *R-Factor*–VoIP transmission quality rating, with a typical range of 50-100. An R-Factor score is derived from multiple VoIP metrics, including latency, jitter, and loss.
- *MOS*–the Mean Opinion Score, which provides a relative measurement of VoIP call quality. PacketShaper measures MOS using a scale of 10-50. To convert to a standard MOS score (which uses a scale of 1-5), divide the PacketShaper MOS value by 10.

These metrics can be measured for real-time transfer protocol (RTP) traffic flows. By configuring your PacketShapers to generate VoIP metrics, you can:

- Assess the suitability of a network to carry VoIP traffic
- Corroborate anecdotal feedback from VoIP users with actual performance statistics

VoIP Metrics Requirements

The VoIP metrics feature has the following requirements and limitations:

- To enable latency, R-Factor, and MOS measurements of VoIP flows between PacketShapers, you must first change each unit's system variable **Enable Latency Calculations** from **off** (default) to **on** (See Enabling VoIP Latency Calculations).
- The statistical reports of VoIP metrics show performance data between PacketShapers, not between VoIP phones.
- The VoIP metrics feature can measure traffic only from VoIP applications whose data is classified as RTP-I. For instance, the VoIP metrics feature does not provide metrics for DialPad, iChat, Vonage, and Skype.
- VoIP latency, R-Factor, and MOS cannot be measured when a PacketShaper is in watch mode.

Enabling VoIP Latency Calculations

When enabled, VoIP latency calculations add a small amount of additional traffic to your network. If you already have a rate policy guaranteeing a minimum amount of bandwidth to each RTP flow, you should increase that amount by three to five percent to avoid a negative impact on call quality. For example, if your VoIP application requires 28k, you should increase its rate policy to 30K when measuring latency. To minimize the impact of VoIP latency calculation overhead, we recommend that you enable this feature only when actively collecting and analyzing data.

To enable VoIP latency calculations, do the following on each PacketShaper between which you want to measure RTP latency:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, select **system variables**.
3. In the **VoIP Metrics** section, ensure that the **Enable Latency Calculations** is set to **on** (default is **off**).
4. Optionally, alter the **Latency probe send interval** to configure the number of seconds between the issuance of VoIP latency probes. PacketShaper uses these probes, which mimic RTP traffic, to compute a VoIP packet's round-trip time between PacketShapers.
5. Click **apply changes**.

> **Note:** Although you must enable latency calculations to measure latency and to compute MOS and R-factor values, you do *not* need to enable latency calculations to measure packet loss and jitter.

Viewing VoIP Metric Statistics

To view reports derived from VoIP metric statistics, see Generate VoIP Reports.

Data for VoIP metrics are stored in measurement variables. (See Measurement Variables for a complete list). The following table describes five useful metrics and their associated measurement variables:

VoIP Metric Type	Measurement Variable
Average Latency (milliseconds) Latency calculations are supported by synthetic RTP packets (probe requests) which are issued by PacketShapers at a configurable interval when an RTP flow is active. The header of the synthetic RTP packet contains a timestamp, sequence number, and unit ID information which are used to measure the round-trip time between the two PacketShapers. For instance, if the sum of measured latency is 500 ms and the number of latency samples is 25, then the average latency is 20 ms. Formula: rtp-latency-sum / rtp-latency-count	rtp-latency-avg
Packet Loss (%) Calculation of packet loss (that is, expected packets that are not received) is performed per RFC 3550 using RTP header sequence numbers. Formula: 100 x (1- packets seen / packets expected), where packets expected = highest RTP sequence number - lowest RTP sequence number +1.	rtp-lost-pkt%
Average Jitter (milliseconds) Jitter, that is, the variation in the delay of received packets in a flow, is measured by comparing the interval when RTP packets were sent to the interval at which they were received. For instance, if packet #1 and packet #2 leave 50 milliseconds apart and arrive 60 milliseconds apart, then the jitter is 10 milliseconds. Formula: rtp-jitter-sum / rtp-jitter-count	rtp-jitter-avg

VoIP Metric Type	Measurement Variable
R-Factor Rating Factor, a numerical score derived from other VoIP metrics, such as latency, jitter, and packet loss, per ITU-T Recommendation G.107. A typical range is 50-90, with a score of 80 or higher indicating satisfactory VoIP call quality.	rfactor
MOS Mean Opinion Score, which provides a relative measurement of VoIP call quality per ITU-T Recommendation G.107. PacketShaper measures MOS using a scale of 10-50. A typical range is 25 to 45, with a score of 40 or higher indicating satisfactory VoIP call quality. **Note**: to convert to a standard MOS score (which uses a scale of 1-5), divide the PacketShaper MOS value by 10.	mos-x10

Response Time Measurement Tasks

PacketShaper's response time measurement (RTM) enables you to measure the quality of users' network experience. RTM provides performance statistics, threshold monitoring, high-level problem indicators, and performance graphs. You can use the response-time analysis to highlight problems and confirm the effectiveness of solutions. For more information, see " RTM Overview " on page 104.

Monitor Response Times

To get an overview of response-time statistics for each class that has current RTM data, you can display the *Monitor Response Time* window. This window indicates each class' total delay threshold, number of transactions, and percent of good transactions, and calculates the average transaction delay and transaction round trip time.

To display the *Monitor Response Time* window:

1. Click the **Monitor** tab. The *Monitor Traffic* window appears.
2. Choose **Response Time** on the **Monitor** drop-down list. The *Monitor Response Time* window appears. **show screen...**

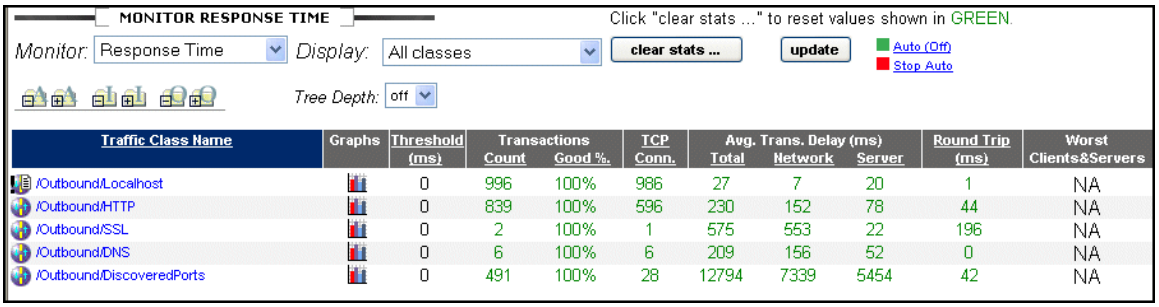

Traffic Class Name	Graphs	Threshold (ms)	Transactions Count	Transactions Good %.	TCP Conn.	Avg. Trans. Delay (ms) Total	Avg. Trans. Delay (ms) Network	Avg. Trans. Delay (ms) Server	Round Trip (ms)	Worst Clients&Servers
/Outbound/Localhost		0	996	100%	986	27	7	20	1	NA
/Outbound/HTTP		0	839	100%	596	230	152	78	44	NA
/Outbound/SSL		0	2	100%	1	575	553	22	196	NA
/Outbound/DNS		0	6	100%	6	209	156	52	0	NA
/Outbound/DiscoveredPorts		0	491	100%	28	12794	7339	5454	42	NA

3. *Optional:* To display only the classes with thresholds set, choose **Classes with Thresholds** on the **Display** drop-down list.

PacketShaper uses the data collected since you last cleared the statistics. Measurement data is cleared with the **clear stats** button in the *Monitor Response Time* window or when you reset the unit or the measurement data.

> **Note:** The *Monitor Response Time* screen lists all classes that have ever had response time data recorded for it. But if a class' matching rules have subsequently been changed in such a way that RTM data is no longer being recorded for the class, PacketShaper will not let you set a threshold.

Understanding the Monitor Response Time Window

The *Monitor Response Time* window contains the following information:

Column	Description
Traffic Class Name	The traffic class being summarized; class names are linked to their response-time windows
Graphs	The icon displays the response-time graphs. See View Response-Time Graphs.

Column	Description
	The acceptable response time for the class. To set the total delay threshold, see Set Thresholds.
Threshold (ms)	If a dash appears in this column, a network or server delay threshold has been set in the command-line interface. (The browser interface only supports *total* network delay thresholds.) Note that using the browser interface to make any changes to a class s RTM settings will clear the network or server delay you set in the CLI.
Transactions Count	The number of transactions for the traffic class since you last cleared the statistics
Transactions Good%	The percent of transactions that completed within the total delay threshold
TCP Conn.	The number of TCP connections established by the class and used by the transactions that were counted in the summary
Avg. Trans. Delay (ms)	The average time in milliseconds the transactions used. The delays are divided into total, network, and server delays. See also Show Transaction Delays.
Round Trip (ms)	An indicator of transit delay that is independent of data size. Round Trip Time (RTT) is the average number of milliseconds spent in transit when a client and server exchange the SYN (synchronize sequence numbers flag) and its corresponding ACK (acknowledge flag). A transaction involving a large amount of data requires the data to be divided into multiple packets. Whereas a transaction's network delay reflects the total transit time for all required packets, the RTT reflects the time for a single small packet to make its way from client to server and another packet to make the return trip. Use the RTT to determine if a large network delay is due to large transactions or a slow network. If the RTT is much smaller than the network delay, the transactions were large. If the two averages are close, a sluggish network caused the longer network delays. Average RTT figures are tracked on a per-host basis. When a traffic class' transactions reference many hosts, the separate figures are averaged.
Worst Clients & Servers	Links to the Worst Clients or Worst Servers lists, described in List the Worst Clients and Servers. If **disabled** appears in this column, a network or server delay threshold has been set in the command-line interface. The Worst Clients & Servers feature works with *total* delay thresholds only.

To reorder the summary's list of classes in ascending or descending order according to a single column's values, click the corresponding column's header.

Show Transaction Delays

To display transaction delays for a class:

1. Click the **Manage** tab. The *Traffic Class* window appears.
2. In the left window pane, select the class.
3. Click **statistics** and choose **response time**. The *Statistics: Response Time* screen appears. **show screen...**

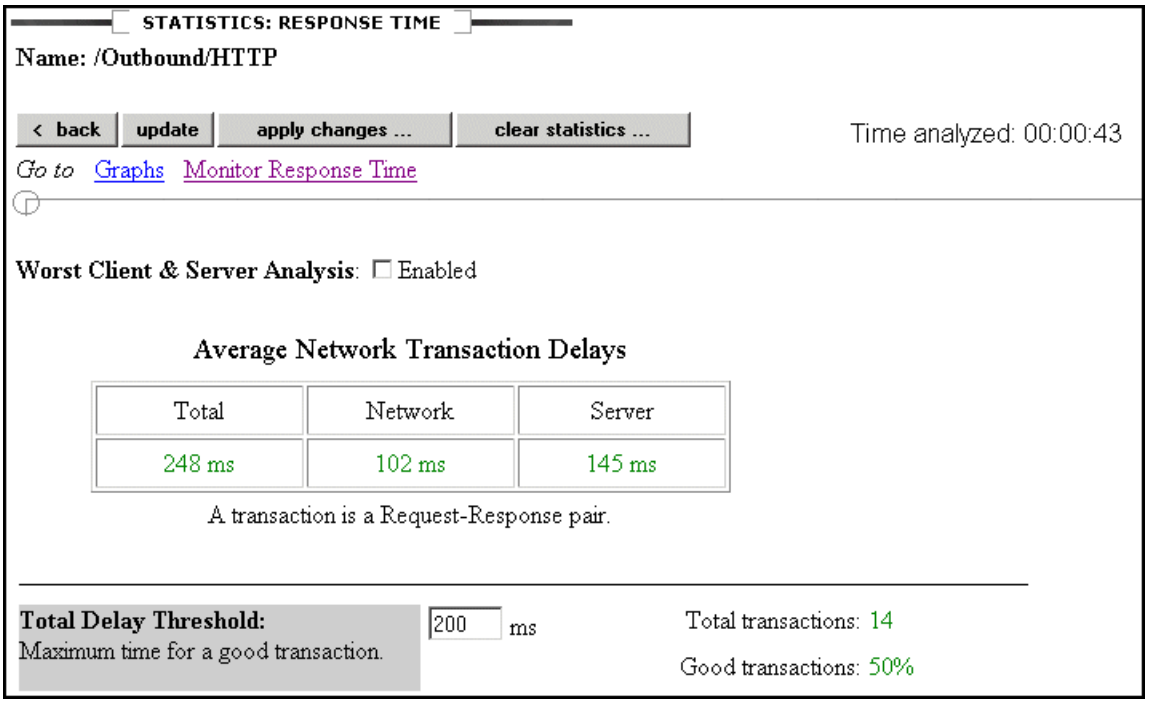

PacketShaper accumulates delay statistics for a running average until you click **clear statistics** at the top of the window. If no average delay figures appear, no traffic has been measured since the statistics were last cleared. Click **update** to refresh the delay averages and total transactions with current information.

Each transaction's response is divided into three delay figures:

Statistic	Description
Total Delay	The number of milliseconds beginning with a client's request and ending upon receipt of the response Transactions that require a large quantity of data to be transferred are divided and sent in multiple packets. This results in longer total and network delays, and does not necessarily reflect a slow network.
Network Delay	The number of milliseconds spent in transit when a client and server exchange data Network delay includes the transit time for all packets required for a request-response transaction. The amount of time the server uses for processing a request is not included in the network delay.

Statistic	Description
Server Delay	The number of milliseconds the server uses to process a client's request after it receives all required data Server delay is the time between when the server receives the last request packet and when it sends the first packet of the reply (the time the server takes to process the client's request).

Note: The network and total delays increase when router queues grow. These two delay figures can be improved by using PacketShaper's bandwidth management features.

Set Thresholds

You can set two types of thresholds:

- Total Delay Threshold
- Service Level Threshold

Total Delay Threshold

The *total delay threshold* allows you to specify acceptable response times for selected traffic classes. Response times are defined as "bad" or "good" according to the threshold you set. When a transaction's total delay is less than the threshold, the transaction is considered good. When the total delay is more than the threshold, the transaction is considered bad.

PacketShaper allows you to have an unlimited number of classes with RTM thresholds.

To set the total delay threshold value for a class:

1. Click the **Manage** tab. The *Traffic Class* window appears.
2. In the left window pane, select the traffic class. The *Traffic Class* window displays the definition for the selected class.
3. Click **statistics** and select **response time**. The *Statistics: Response Time* screen appears. **show screen...**

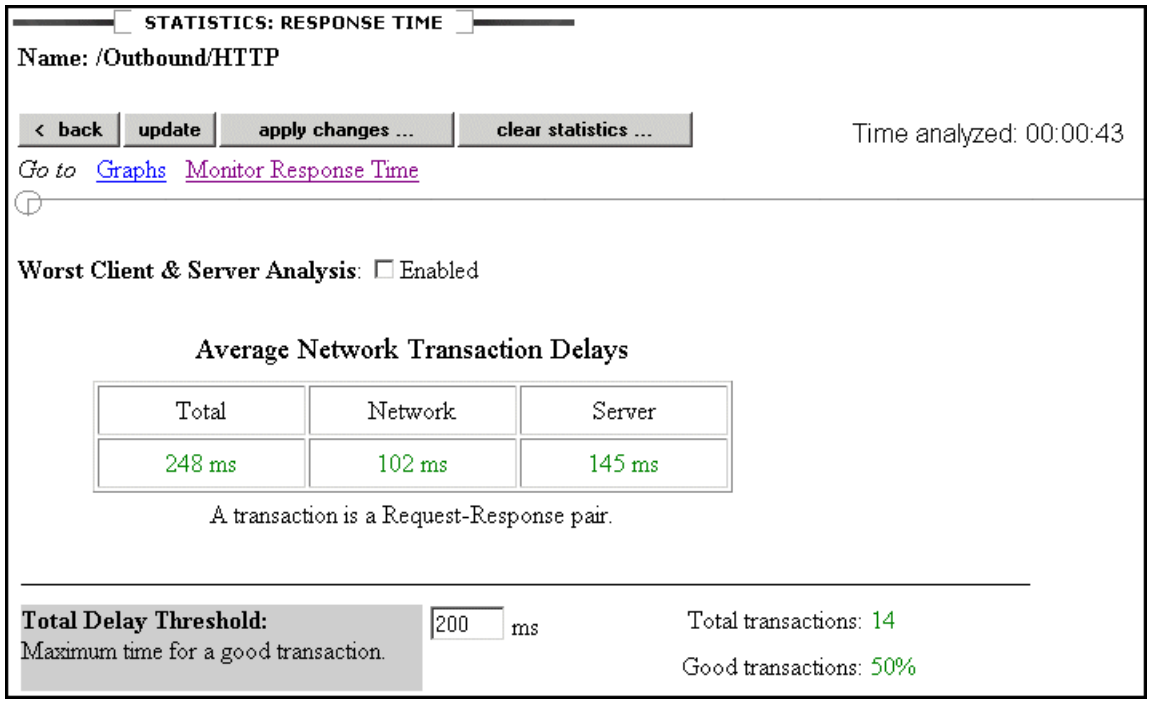

If **response time** is not listed on the **statistics** menu, PacketShaper does not have any RTM data for the selected class

4. Enter the desired threshold value in the **Total Delay Threshold** field. To determine the number of milliseconds that constitute "too slow" for a total delay, use one of following methods:

- Use the value PacketShaper suggests.
- Determine the current average total delay for the traffic class. To determine the current average total delay, compare the average total delay with users' perception of performance. Increase this number until it represents an unacceptable delay (not a typical delay).

5. Click **apply changes**.

The accumulated delay statistics for the class are deleted, and PacketShaper begins timing new transactions and calculating new running delay averages. In addition, the **Service Level Threshold** field is now available. (See Service Level Threshold.)

Once you have set a threshold, the *Statistics: Response Time* window displays the percentage of good transactions. This enables you to quickly evaluate the percentage of transactions that are within an acceptable limit of the required service level.

Service Level Threshold

After setting the total delay threshold, use the *service level threshold* to specify an acceptable percentage of transactions completed within one-minute intervals.

> **Note:** The Service Level Threshold field does not display until you have set the total delay threshold.

PacketShaper tracks the percentage of transactions during each one-minute interval that complete within the number of milliseconds indicated in the total delay threshold. If the percentage is less than the specified service level threshold, that interval is considered unacceptable. PacketShaper indicates the number of unacceptable intervals and the time of the last unacceptable interval in the *Statistics: Response Time* window.

To set the service level threshold:

1. In the **Service Level Threshold** field, enter a percentage that represents the minimum level of good transactions that is acceptable for that traffic class. For example, if you enter 90 for the service level threshold, then 90 percent of transactions in the class must be good (that is, below the total delay threshold). **show screen...**

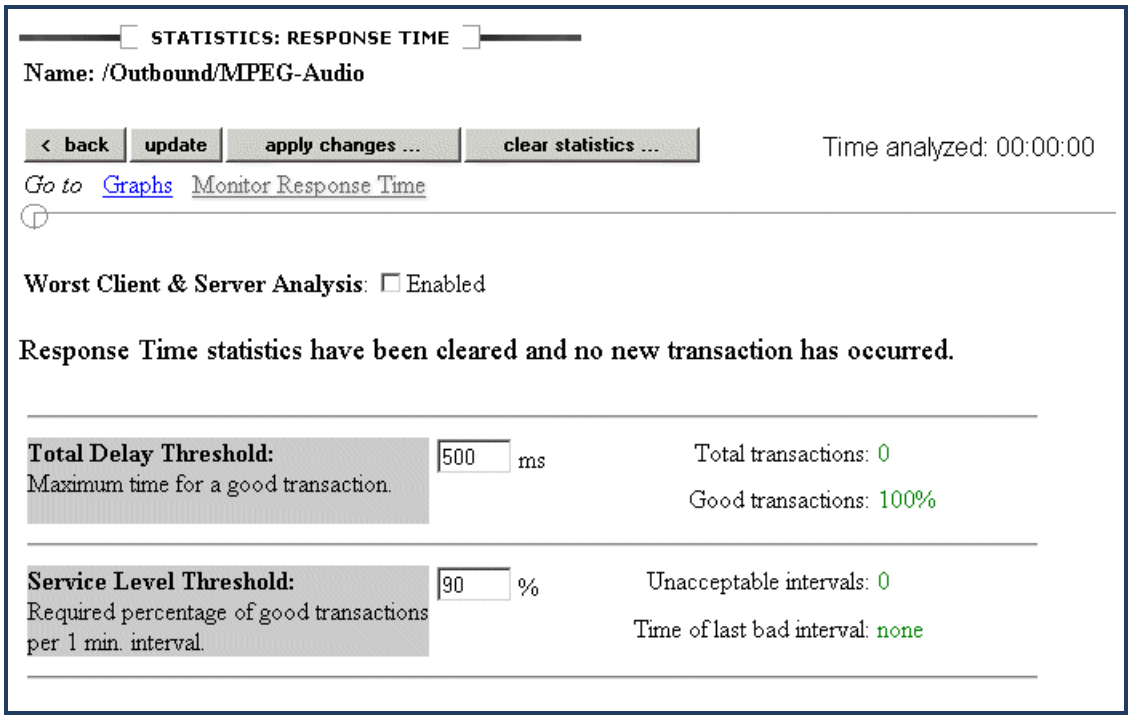

> **Tip:** Use a number that will prompt a corrective action or investigation if the percentage falls below the acceptable level.

2. Click **apply changes**.

Tracking Service Level Agreements

PacketShaper's two performance thresholds (the total delay threshold and the service level threshold) enable you to quantify and track Service Level Agreements.

For example, say you promise your users that 90 percent of their HTTP transactions in each one-minute interval will complete in under 600 ms. To track this, set the total delay threshold to define a good transaction as being under 600 ms. Then, use the service level threshold to define acceptable performance over time as being 90 percent good transactions. You may want to create an event to automatically notify you if the percentage of good transactions drops under the service level threshold.

View Response-Time Graphs

To view response-time graphs of a class:

1. Click the **Monitor** tab. The *Monitor Traffic* window appears.
2. Choose **Response Time** on the **Monitor** drop-down list. The *Monitor Response Time* window displays. **show screen...**

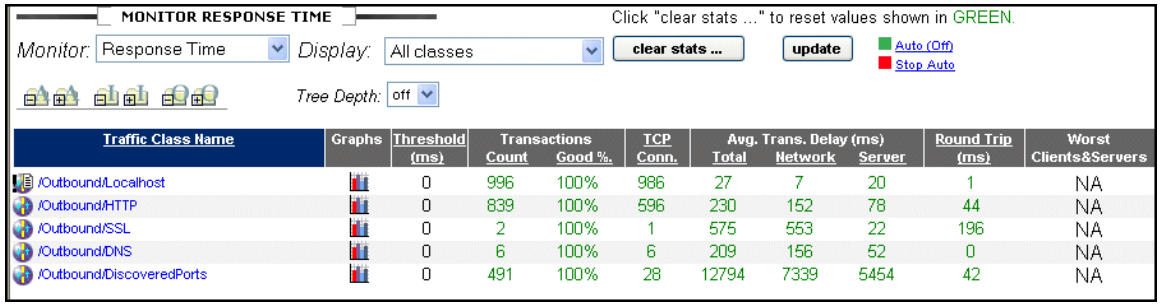

Traffic Class Name	Graphs	Threshold (ms)	Transactions Count	Transactions Good %	TCP Conn.	Avg. Trans. Delay (ms) Total	Avg. Trans. Delay (ms) Network	Avg. Trans. Delay (ms) Server	Round Trip (ms)	Worst Clients&Servers
/Outbound/Localhost		0	996	100%	986	27	7	20	1	NA
/Outbound/HTTP		0	839	100%	596	230	152	78	44	NA
/Outbound/SSL		0	2	100%	1	575	553	22	196	NA
/Outbound/DNS		0	6	100%	6	209	156	52	0	NA
/Outbound/DiscoveredPorts		0	491	100%	28	12794	7339	5454	42	NA

3. For the class you want to graph, click the icon. Three graphs appear in the window. **show screen...**

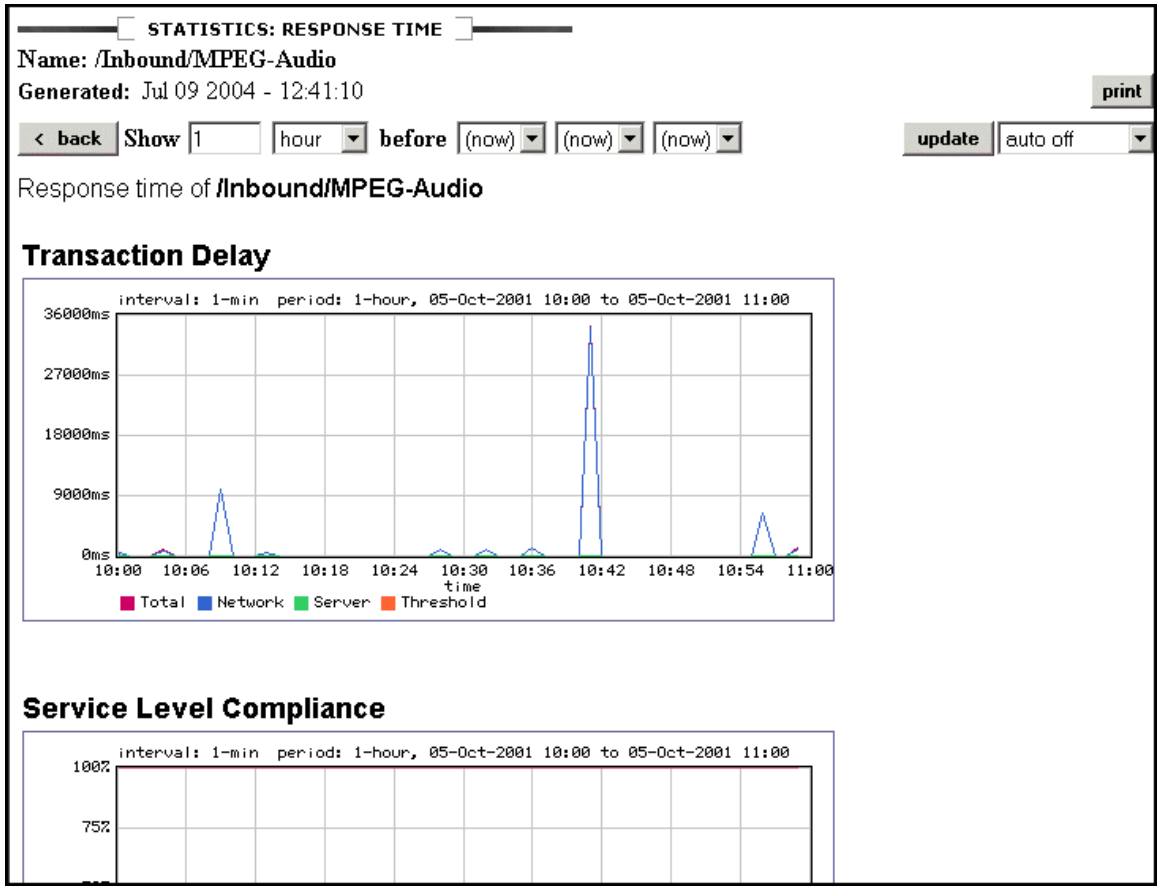

Graph	Description
Transaction Delay	A timeline of a traffic class' average response times. Three colored lines track the total, network, and server delays, so you can determine the source of any significant delays. The horizontal line represents the Total Delay Threshold as a reference point to determine the quality of the response times.
Service Level Compliance	The percentage of good transactions. This graph shows two lines, one representing the required percentage of good transactions and the other representing the actual percentage of good transactions. The first line is always flat. The second line varies, depending on performance.
Transaction Delay Distribution	The number of transactions whose total delay falls into each of the 14 transaction-time buckets; the median delay is indicated on the bottom of the graph. Each time slot contains the number of transactions that are delayed with the specified number of seconds.

4. Click **More details** at the bottom of the window to display a report containing three additional graphs.

Graph	Description
Network Delay Distribution	The number of transactions whose network delay falls into each of the 14 transaction-time buckets; the median delay is indicated on the bottom of the graph.
Server Delay Distribution	The number of transactions whose server delay falls into each of the 14 transaction-time buckets; the median delay is indicated on the bottom of the graph.
Packet Round-Trip Time	The history of a transaction's round-trip time (RTT). RTT is the average number of milliseconds spent in transit when a client and server exchange the SYN (synchronize sequence numbers flag) and its corresponding ACK (acknowledge flag). A transaction involving a large amount of data requires the data to be divided into multiple packets. Whereas a transaction's network delay reflects the total transit time for all required packets, the RTT reflects the time for a single packet to make its way from client to server and another packet to make the return trip.

5. To graph a different time period, see Change the Time Period.
6. Click **print** to print the report.

Enable Worst Client and Server Analysis

For traffic classes with total delay thresholds, PacketShaper maintains lists of the "Worst Clients" and "Worst Servers." Each list displays up to 20 clients or servers that have the highest percentage of transactions exceeding the total delay threshold. These lists are useful for determining the scope of a problem and for focusing a diagnostic search. If you want PacketShaper to track the worst performers in a class, you must first enable worst client and server analysis.

Note: You cannot enable worst client and server analysis without having a total delay threshold. You can either set the threshold first and then enable worst analysis or set the threshold at the same time as you enable worst analysis.

Caution: If you set a network or server delay threshold via the command-line interface, you should not enable worst client and server analysis. Doing so will clear the threshold setting. (Worst client and server analysis is for *total* delay threshold only.)

To enable worst client and server analysis:

1. Click the **Manage** tab. The *Traffic Class* window appears.
2. Select the traffic class. The *Traffic Class* window displays the definition for the selected class.
3. Click **statistics** and select **response time**. The *Statistics: Response Time* screen appears. **show screen...**

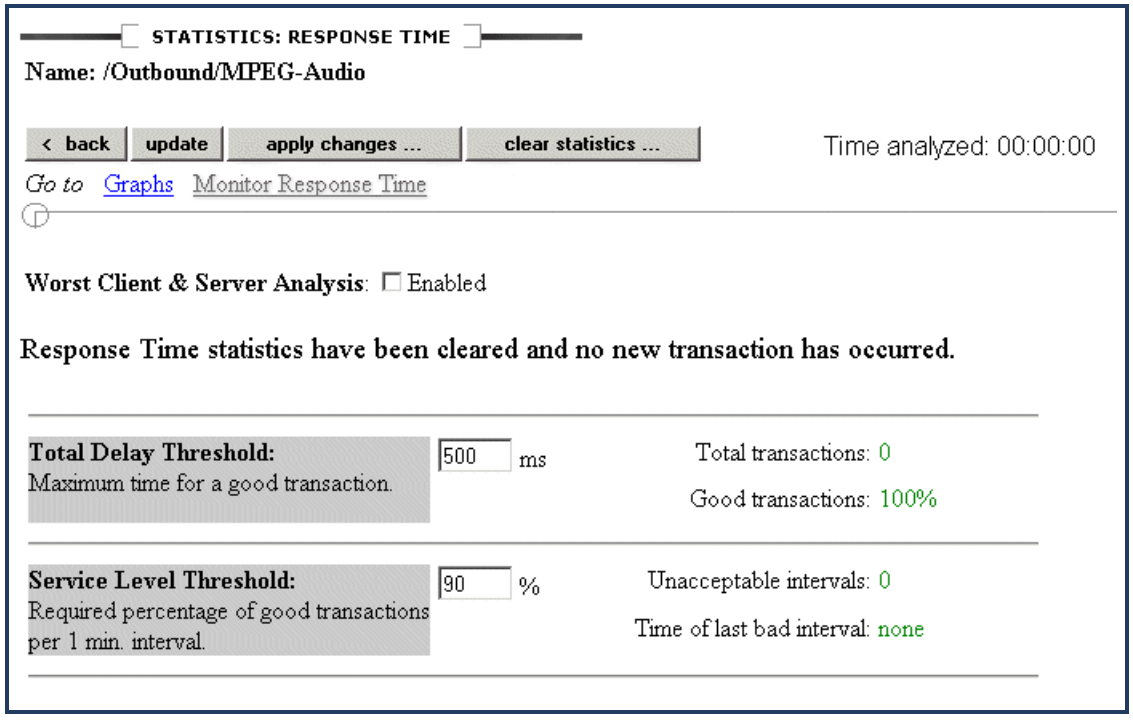

4. Select the **Worst Clients & Servers Analysis** checkbox.
5. Click **apply changes**.

Links to **Worst Client** and **Worst Server** appear near the top of the window. **show screen...**

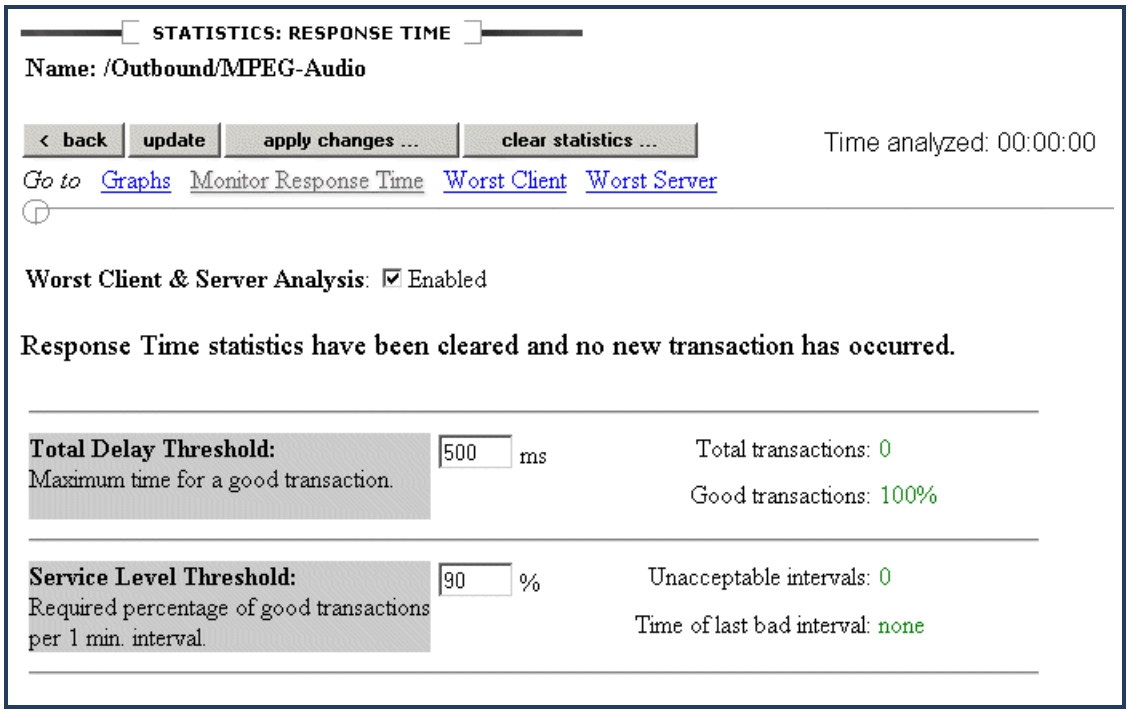

You can click these links to display the associated lists.

List the Worst Clients and Servers

In order to list the worst clients and servers for a class, two conditions must be met: you must have set thresholds and you must have enabled worst client and server analysis on the class.

To display a list of the hosts with the worst response times in a class:

1. Click the **Manage** tab. The *Traffic Class* window appears.
2. In the left window pane, select the class.
3. Click **statistics** and choose **response time**. The *Statistics: Response Time* screen displays. **show screen...**

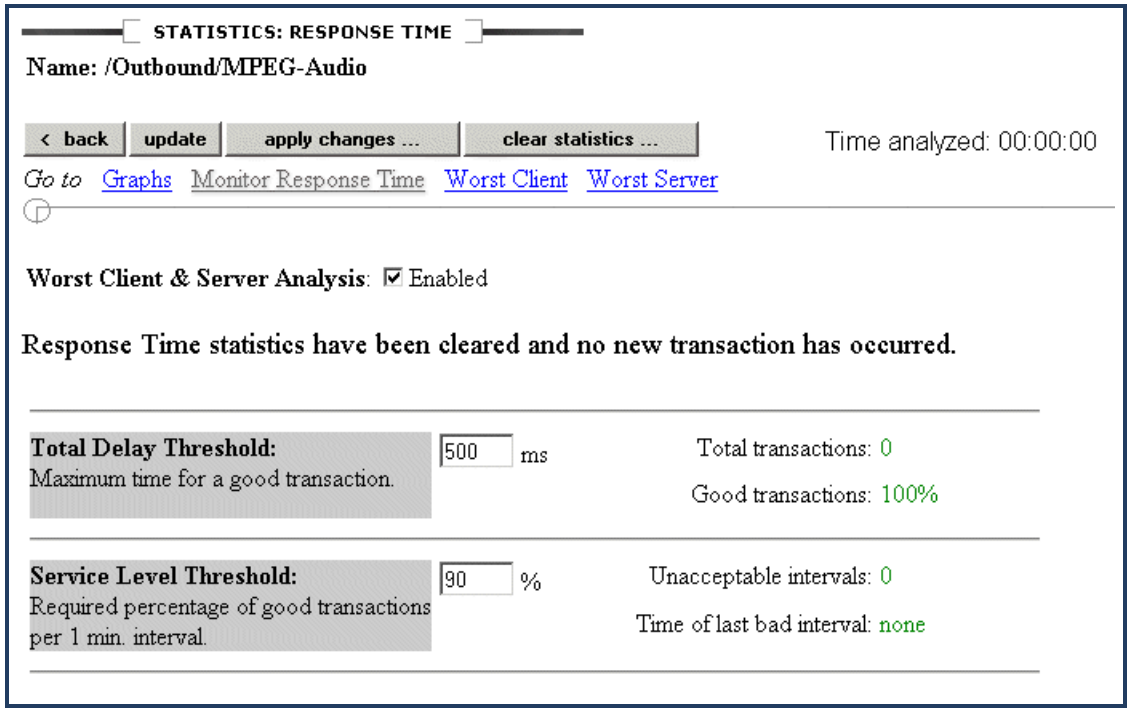

4. Click **Worst Client** or **Worst Server** at the top of the screen. The *Worst Clients* or *Worst Servers* screen appears.

show screen...

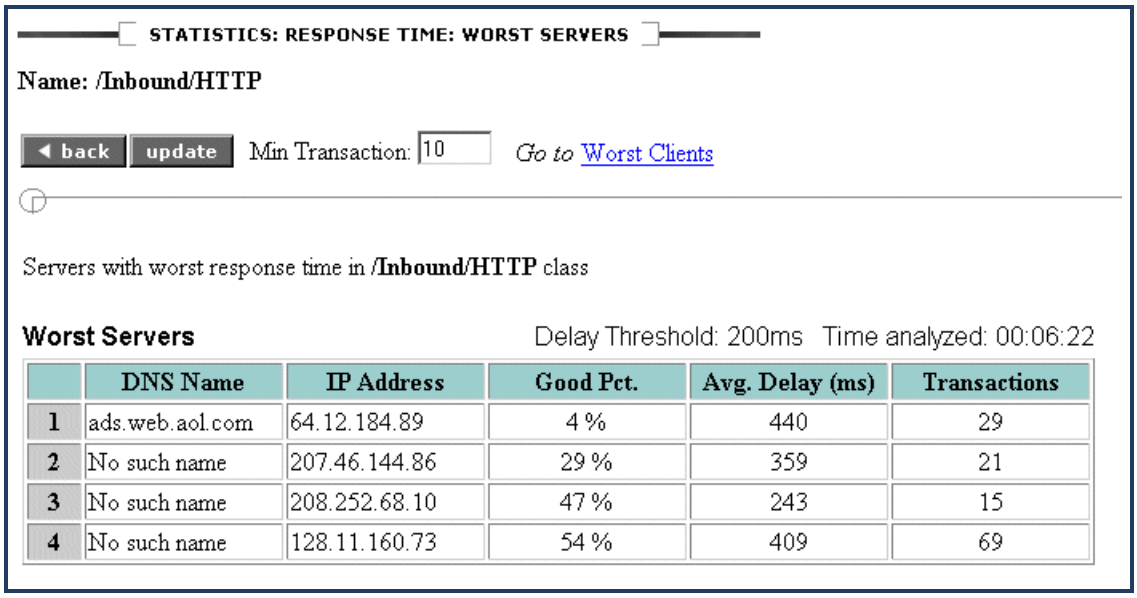

	DNS Name	IP Address	Good Pct.	Avg. Delay (ms)	Transactions
1	ads.web.aol.com	64.12.184.89	4 %	440	29
2	No such name	207.46.144.86	29 %	359	21
3	No such name	208.252.68.10	47 %	243	15
4	No such name	128.11.160.73	54 %	409	69

The *Worst Clients* and *Worst Servers* screens display the DNS Name, IP Address, and statistics for the hosts with the worst response times. The following statistics are displayed for each "worst" host:

- **Good Pct.** The percentage of transactions whose delays are less than the specified total delay threshold. (The *Delay Threshold* is also indicated on this screen, above the list.)
- **Avg. Delay (ms)** The average transaction delay in milliseconds; includes network and server delays.
- **Transactions** The total number of transactions since you enabled Worst Client & Server Analysis. (The *Time analyzed* the total number of hours and days appears on the screen, above the list.)

The information displayed on the screen is captured at the moment you click the **Worst Client** or **Worst Server** link.

- Use **update** to refresh the screen with the most current data.
- Use the **Min Transaction** field to specify the minimum number of transactions a host must have in order to be listed on this screen. By default, hosts are not listed on this screen unless there has been at least 10 transactions. After entering a new value, click **update**.

Understanding "Worst"

When viewing the Worst Servers and Worst Clients lists, keep in mind that this list is based on a percentage. In the following example, think about which server is considered to be the worst server.

Example of Worst Server

The Total Delay Threshold is set to 3 milliseconds.

- Server A has transactions with the following times in milliseconds: 1, 1, 2, 1, 1000
- Server B has transactions with these times: 1, 1, 4, 4, 1

Server B is the Worst Server because it has the highest percentage of transactions (40% versus 20% for Server A) exceeding the Total Delay Threshold.

Adaptive Response Tasks

The adaptive response feature makes it easier to monitor your network by enabling an individual PacketShaper to monitor unit, application, and network health, hosts, traffic classes, links and partitions. When this feature is enabled, the unit displays color-coded status indicators that let you identify potential problems at a glance. You can configure action files to customize and automate responses to problems, and view drill-down incident reports to see exactly what happened when a problem occurred.

This section includes tasks for enabling this feature, creating and editing agents, creating action files, and viewing incident reports.

Enable Adaptive Response

The adaptive response (AR) feature is not turned on by default. To activate this feature:

1. Click the **Setup** tab.
2. On the *basic settings* setup page, turn the **Adaptive Response** setting to **on**.
3. Click **apply changes**.

Note that after you initially enable adaptive response, all the status indicators on the Info tab will be gray. The adaptive response dashboard will not display the red, green, or yellow indicators until:

- The agents are enabled. (All preconfigured agents are disabled by default.)
- The agents have had time to gather and analyze data (that is, when the agent has reported a score for the first interval). Some of the categories may not have any default agents, so these indicators will remain gray unless you create an agent in that category.

Enable Adaptive Response Agents

The preconfigured agents are created on your PacketShaper or PolicyCenter configuration when you first enable the adaptive response feature. All preconfigured agents are disabled by default, and you need to decide which ones you want to use and enable the ones that look most useful. Once an agent is enabled, adaptive response will begin monitoring the variables and reporting scores associated with the agent.

To enable AR agents:

1. Click the **Setup** tab.
2. Choose **adaptive response** from the **Choose Setup Page** list.
3. For each agent you want to enable, click the check box in the **Agent On** column.
4. Click **apply changes**.

> **Note:** The preconfigured agents are only created when there are no custom or predefined agents currently on your PacketShaper or PolicyCenter configuration. If you delete a preconfigured agent, toggling the adaptive response setting off and then on again will not recreate that agent unless all agents are also removed.

Add a New Adaptive Response Agent

Each PacketShaper or PolicyCenter configuration can have a maximum of 32 agents. Each new agent you create is based on one of the existing agent templates.

To create a new agent on a PacketShaper or a PolicyCenter configuration:

1. Click the **Setup** tab.
2. From the **Choose Setup Page** list, choose the **Adaptive Response** setup page.
3. Click **add**. The *add agent entry* window opens.
4. Click the **create a new agent from this agent template** drop-down list, and choose the type of agent you want to create.
5. Enter a unique identifying name for the new agent in the **Name** field. An agent name can have up to 32 alphanumeric characters, including -, _, and . (period).
6. Specify an evaluation interval, in minutes (how often the agent should measure its target).
7. If you have already created a red or green action file, you may specify the filenames in the **Red Action File** and **Green Action File** fields. You must include the .cmd file extension with the filename. If you have not yet created action files, you may do so later.
8. Click **OK and Edit Parms**. This opens the *Edit Agent Entry* window where you can define the parameter values for your new agent.
9. The Agent Parameters section lists all scoring parameters for the agent, as well as any user-defined variables in the green or red action files. Enter a new value for any parameter you wish to modify. The definition of each field depends on the template you selected. For descriptions of the parameters associated with each template, see Adaptive Response Agents Overview and select the agent template name.
10. Click **OK** to save your changes.

Edit an Adaptive Response Agent

Once you have created an agent on your unit, you can edit its thresholds and parameters as needed.

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **adaptive response**.

3. Click the **edit** button beside the agent you want to modify.

4. *For PacketShapers assigned to a PolicyCenter configuration:* Agents inherited from a parent configuration can be viewed but cannot be edited unless they are overridden with a local copy of that agent. To override an inherited agent, click the **inherited** check box (remove the check mark) beside the agent you want to edit, then click **apply changes**. The **view** button by the agent name becomes an **edit** button, allowing you to edit the agent on that child configuration. The agent in the parent configuration will not change.

5. The Agent Parameters section lists all scoring parameters for the agent, as well as any user-defined variables in the green or red action files. Enter a new value for any parameter you wish to modify. The definition of each field depends on the template you selected. For descriptions of the parameters associated with each template, see Adaptive Response Agents Overview and select the agent template name.

6. Click **OK** to save your changes.

Delete an Adaptive Response Agent

Use the following procedure to delete an existing agent from a PacketShaper or a PolicyCenter configuration:

1. Click the **Setup** tab.

2. Choose **adaptive response** from the **Choose Setup Page** list.

3. Click the **delete** check box beside the agent you want to delete, then click **apply changes**.

> **Note:** To restore the default set of agents, use the agent createdefaults CLI command. You will not be able to permanently delete an agent from a PolicyCenter configuration if that configuration inherits the agent from its parent. Adaptive response agents can only be deleted from the configuration on which they were created.

Temporarily Disable an Agent

If you do not want to permanently delete an agent, you may instead turn the agent off. A disabled agent no longer measures its target or reports values, but the agent still retains its current values and settings.

To temporarily disable an agent:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **adaptive response**, or click the **Adaptive Response Setup** link.

3. Click the check boxes in the **Agent On** column (removing the check mark) beside the agent you want to disable, then click **apply changes**.

To reenable an agent that has been turned off,

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list,select **adaptive response**.

3. Click the check boxes in the **Agent On** column (adding the check mark) beside the agent you want to reenable, then click **apply changes**.

Add or Edit Adaptive Response Action Files

The adaptive response feature allows you to set thresholds for acceptable and unacceptable performance, and specify command-file action scripts to run when these thresholds are reached. These *action files* of CLI commands can make automatic changes, often correcting problems before users are ever impacted. The action file scripts can also be configured to send notifications via email, syslog, and SNMP traps. Once the agent status returns to normal, a second action file can undo the changes, or run additional commands.

The steps for creating and editing action files differ depending on whether you intend them to run on a unit in local (standalone) mode, or a unit in shared mode assigned to a PolicyCenter configuration. A PacketShaper in local mode requires just one instance of each action file, and so its action files can be created and edited directly on the unit. When a PacketShaper is in shared mode, its action files must reside on both the unit's data disk *and* the PolicyCenter server. PolicyCenter action files should be created on the PolicyCenter server first, and then distributed to PolicyCenter configurations and units.

PolicyCenter Action File Portfolios

PolicyCenter also requires that action files be grouped into *portfolios*, folders containing one or more action files that can be prescribed to a PolicyCenter configuration and its assigned units. If you are creating your first PolicyCenter action file, or a new portfolio of files, you may need to create a new Symantec Systems/PolicyCenter/publish/agent/*<portfolio_name>* folder on your PolicyCenter server with the Windows Explorer utility. Create a directory structure with one or more portfolios of action files (for example, *Blue Coat Systems/PolicyCenter/publish/agent/actndir1* and *Blue Coat Systems/PolicyCenter/publish/agent/actndir2*) that each contain one or more action files. A portfolio name can be up to eight characters long, including 0-9, a-z, A-Z, -,_, and . (period). Spaces are not allowed in a portfolio name, and you will not be able to prescribe a portfolio unless that portfolio folder contains at least one action file.

Every PolicyCenter action file you create should have a unique name, regardless of what portfolio it is in. (For example, you shouldn't create the file *../publish/agent/prtfolo1/action.cmd* if there is an existing file *../publish/agent/prtfolo2/action.cmd*.) If you give a new action file the same name as an existing action file, prescribing the new action file will overwrite the existing file on both the unit and any PolicyCenter configuration that uses that file.

> **Note:** If you wish to distribute action files that are currently on a PacketShaper but not yet on the PolicyCenter server, FTP the action files from the PacketShaper directory *9.258/agent/cmd* into the PolicyCenter directory ../PolicyCenter/publish/agent/*<portfolio_name>*, then use the action file distribution feature.

Creating a New Action File

Action files must be stored in your PacketShaper's *9.258/agent/cmd* directory or if you are using PolicyCenter, in the Blue Coat Systems/PolicyCenter/publish/agent/*<portfolio_name>* directory. PacketShaper's File Browser utility has a facility for creating command files. Use one of the following ways to access the File Browser:

- The **file browser** button on the Info tab
- The **browse files** button on the *Edit Agent* page

The second method offers two advantages. First, it automatically stores the action file in the proper place. (You don't have to navigate to the directory). Second, after creating the action file, you are on the right screen for supplying the red or green action filename. The steps below use this method.

To create a new action file:

1. Click the **Setup** tab and click the **Adaptive Response Setup** link.

2. Click the **edit** button next to the agent for which you want to create the action file.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored). In PolicyCenter, navigate to the desired portfolio (Blue Coat Systems/PolicyCenter/publish/agent/*<portfolio_name>*).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, AZ, -,_, and . (period). Spaces are not allowed. Example: *red-def.cmd*

6. In the Contents area, enter the CLI commands for your action file. See Action Files Overview for lists of variables you can include.

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** or **Green Action File** field. You must include the .cmd file extension.

9. Click **OK**.

10. *(For PolicyCenter only)* You must now prescribe the portfolio with the new action file to a PolicyCenter configuration. See "Distribute Action Files" In PolicyCenter PacketGuide for details on distributing action files to PolicyCenter configurations and units in shared mode.

Action Output Files

After action files have run, an output file is created in the unit directory *9.258/agent/cmd/complete*. The output file has the same name as the action file, but with a numbered file extension (for example, *red-host.17*). Here's an example output file:

```
*** (9.258/AGENT/CMD/pending/testpart.61, line 1): 'send email raltman@test.com "FTP
Partition" "The usage for inbound/ftp exceeded 90% of the partition size"'
```

Editing an Action File

You can edit an action file directly on an individual unit in local mode, or, for action files for a PolicyCenter configuration, from the PolicyCenter browser interface. Remember, units in shared mode require that its action files reside on both the unit's data

disk *and* the PolicyCenter server. After you edit a file on the PolicyCenter server, you must use the file distribution feature to resend the updated file to the units assigned to that configuration.

1. *To edit an action file on a unit in local mode:* Click the **Info** tab, and click the **file browser** button. Navigate to the folder 9.258/agent/cmd.

-or-

To edit an action file for a unit in shared mode using a PolicyCenter configuration: Access the PolicyCenter browser interface, click the **file browser** button on the **Info** tab, and navigate to the folder Blue Coat Systems/PolicyCenter/publish/agent/*<portfolio_name>*.

2. Click the **edit** button by the name of the action file you want to change. The action file will open in a separate browser window.

3. Make the desired changes, then click **save**.

4. *(for PolicyCenter only)* Use the file distribution feature to distribute the updated action file from the ../publish/agent/*<portfolio_name>* directory to individual PacketShapers and their PolicyCenter configurations.

Another way to edit an action file is to access the File Browser utility when editing the agent. This technique is explained in the Creating a New Action File section.

Reset Action File Parameters on an Agent

When you add user-defined variables to an action file, you must reset the parameters for the agent in order to get the new variables to appear on the *Edit Agent Entry* page. For example, suppose you include the following CLI command in a red action file:

```
send email $EmailAddress $Subject $Message_to_admin
```

These variables (EmailAddress, Subject, and Message_to_admin) will not appear on the *Edit Agent Entry* screen until you reset the parameters.

1. *To reset an action file on a unit in local mode:* Access the unit's browser interface.

-or-

To reset an action file for a unit in shared mode using a PolicyCenter configuration: Access the PolicyCenter browser interface. Click the **Configurations** tab, and select a unit or draft configuration with an agent that references the updated action file.

2. Click the **Setup** tab, and choose **Adaptive Response Settings**.

4. From the *Adaptive Response* page, click the **edit** button beside the agent with the changed action file.

5. Click the **Reset parameters from the (red/green) action file** check box. A pop-up window will appear, warning that resetting the action file parameters will clear all of their values and reload the parameter list. Click **OK** on the pop-up, then click **OK** on the *Edit Agent Entry* window.

6. The new action file parameters have been recognized, but because all parameter values were cleared in the process, you must now reassign values to the user-defined parameters in that file. Click the **edit** button beside the agent name again to reopen the *edit agent entry* window.

7. The variables on the reset red or green action file will display the value *Not Set*. Enter the desired values for these variables, then click **OK**.

Enable or Disable Action Files

If you do not want to permanently delete an agent, you may instead disable action files from the *Adaptive Response* page. This stops the agent action files from running, while still keeping the action file parameters of the agent itself. You may reenable an action file at any time.

To disable an action file from a PacketShaper or a PolicyCenter configuration:

1. Click the **Setup** tab.
2. Choose **adaptive response** from the **Choose Setup Page** list.
3. In the **Red Action Enabled** or **Green Action Enabled** column, click the check box (removing the check mark) beside the action file you want to disable.
4. Click **apply changes**.

To reenable an action file from a PacketShaper or a PolicyCenter configuration:

1. *For PolicyCenter only,* click the **Configurations** tab and choose the configuration with the action file you want to reenable..
2. Click the **Setup** tab.
3. From the **Choose Setup Page** list, select **adaptive response**.
4. In the **Red Action Enabled** or **Green Action Enabled** column, click the check box (adding a check mark) beside the action file you want to reenable.
5. Click **apply changes**.

Distribute Action Files

An *action file* is a text file associated with a specific adaptive response agent that contains a set of PacketShaper CLI commands; when a threshold is crossed for a particular agent, the action file automatically runs these commands. When a PacketShaper unit is in shared mode and subscribed to PolicyCenter, you can have PolicyCenter's distribution server distribute a portfolio of adaptive response action files to PacketShaper units. You will use PolicyCenterto configure the distribution service (set a schedule for file distribution and select the action files to distribute to units in the configuration.) See PacketGuide for PolicyCenter for details.

The *Action File Distribution* setup page on the PacketShaper unit displays the settings that were defined in PolicyCenter: the synchronization schedule and prescribed action files. You can override these settings for the unit, if you need to do so.

View Drill-Down Incident Reports

An incident report is a table or graph that adaptive response automatically generates when an agent changes status. A red incident report shows supporting data for the interval in which the agent crossed the red threshold (the time that the problem occurred). A green incident report shows the data for the interval in which the agent returned to a green status.

Adaptive response generates an incident report if the agent's very first evaluation interval records a red status or when the status changes from:

- green to red
- yellow to red
- blue to red
- red to green
- yellow to green

The actual incident files reside in the directory 9.258/agent/cmd/complete. Adaptive response stores the last generated incident report for each agent.

Agents based on the following agent templates do not generate green incident reports since they aren't applicable:

- High Bandwidth New App
- New Application
- Traffic Performance
- Partition Utilization

View an Incident Report

You can view incident reports (also known as drill-down reports) from the browser interface of an individual PacketShaper, or for units in shared mode, from the PolicyCenter browser interface.

To view an incident report from the browser interface of your unit:

1. Enter the IP address of the unit into your web browser address bar, or from the PolicyCenter Configurations tab, click the IP address of the unit. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each.

2. To view a pop-up window with the current value for each individual agent in a category, hover your mouse over a colored status indicator, as shown below.

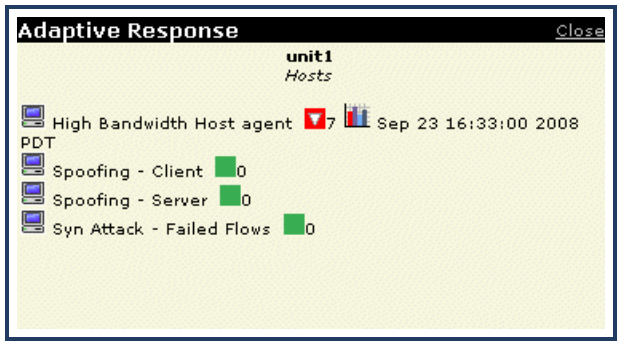

3. If the agent has ever changed status (from green/blue/yellow to red, or from red/yellow to green, a report icon with a link to a detailed report will be available for that agent. Click the report icon to view the report in a separate browser window. For example:

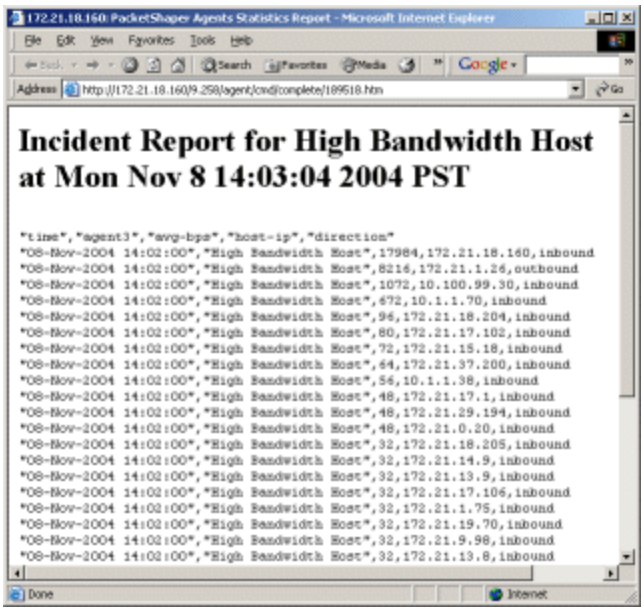

Adaptive Response Agent Examples

Here are examples of how to use the various adaptive response agents:

Network Health

Partition Utilization

Traffic Performance

Unit Health

Memory Allocation

System Load

Unit Limits

Application Health

Default Traffic

New Application

High Bandwidth New App

Hosts

High Bandwidth Host

Spoofing

SYN Attack

User Event Emulation

Partition ME Variables

Class ME Variables

Link ME Variables

Create Partition Utilization Agents

A dynamic way to keep tabs on your partitions is to create a Partition Utilization agent to monitor the utilization of the partitions you have created. When a class' utilization reaches a certain percentage of the partition's maximum size, the adaptive response feature will alert you that a threshold has been exceeded. The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when the class' utilization crosses a predefined threshold; this method requires that you create an *action file*– a text file of PacketShaper commands. The action file can take other actions, as well, such as modifying the partition size.

You can monitor the partition utilization of your Inbound and Outbound links, a specific traffic class, or all classes.

In order to use this feature, the adaptive response feature must be enabled and you must create an agent based on the Partition Utilization template.

Create a New Agent

To create a partition utilization agent that monitors the utilization of one or more partitions:

1. On the **Setup** tab, select the **Adaptive Response** setup page.

2. Click **add**. The *add agent entry* window opens.

3. Click the **create a new agent from this agent template** drop-down list, and choose *Partition Utilization*.

4. Enter a unique identifying name for the new agent in the **Name** field, for example *FTP Partition Utilization*.

5. If desired, change the evaluation interval, in minutes (how often the agent should measure its target).

6. Click **OK and Edit Parms**. This opens the *Edit Agent Entry* window where you can define the parameter values for your new agent.

7. In the **ClassName** field, enter the name of the class to be monitored (for example, */Inbound/FTP*). Note that the class must already have a partition defined. If desired, a wildcard may be used to monitor many classes (for example, */inbound** to monitor all classes in the inbound direction or * to monitor all classes).

8. If desired, modify the threshold values:

   Utilization The percent utilization of the partition which, when exceeded, will cause a class to be in violation. For instance, if the Utilization is set to 90 and the utilization of one of the monitored classes exceeds 90 percent of the partition size, that class will be in violation. Assuming the RedThreshold is set to 1 (the default), the agent's status will change to red.

   RedThreshold The value at which the status of the agent turns red, and a red action file is triggered, if present. For a

Partition Utilization agent, this number represents an unacceptable number of violating classes; a class is in violation if it exceeds the defined Utilization percentage. For instance, if the RedThreshold is set to 1, the agent's status will change to red when at least one class had utilization greater than the defined Utilization percent.

GreenThreshold An acceptable number of violating classes (usually 0).

9. Click **OK** to save your changes.

After creating the agent, you will want to monitor it to see whether usage in the class(es) has come close to the partition's maximum size.

Check the Agent Status

To check the status of the Partition Utilization agent:

1. Click the **Info** tab. The PacketShaper's Info tab shows an icon for each agent category, and a colored status indicator for each. The Partition Utilization agent is in the Network Health category (the first one).

2. To see the status of each individual agent in the Network Health category, hover your mouse over the first colored status indicator, as shown below.

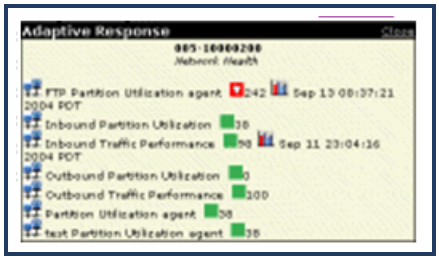

3. In the pop-up window, locate the Partition Utilization agent that you defined. What color is its status indicator?

Green During the last evaluation interval, usage did not exceed the Utilization percentage you defined. In other words, the partition size is appropriate at this time.
Red During the last evaluation interval, usage exceeded the defined Utilization percentage. In other words, usage is close to the partition's maximum size. In this case, you may want to consider increasing the size of the partition (assuming it is traffic you are trying to *protect*, not *contain*). If the status indicator is red, you should look at the incident report for details.

If the agent has ever changed status (from green/yellow to red or from red/yellow to green, a report icon with a link to a detailed report will be available for that agent.

4. If there is a report for the Partition Utilization agent, click the report icon. The Partition Utilization report appears in a separate browser window and lists all the classes during the interval whose utilization was higher than the specified

Utilization percentage.

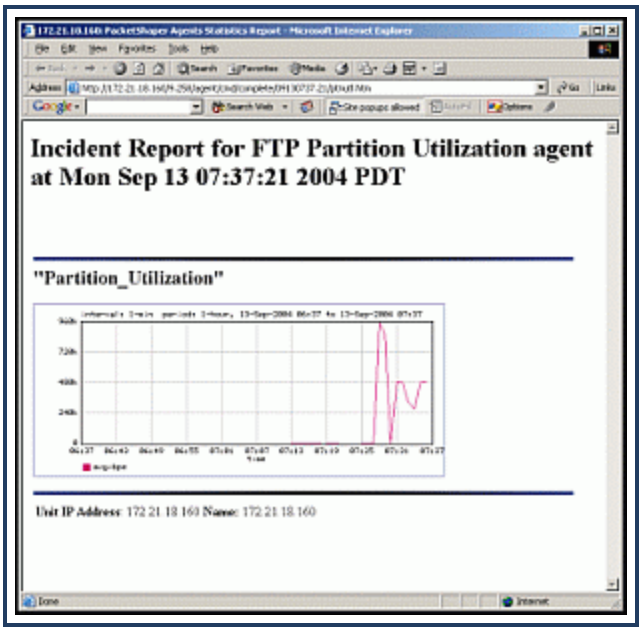

Create an Action File that Sends an Email Notification

If you would like to be notified by email when the class' usage is approaching the partition's maximum size, you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file that sends an email notification:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and select the **Adaptive Response** page.

2. Click the **edit** button next to the partition agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, AZ, -,_, and . (period). Spaces are not allowed. Example: *ftp-part*

6. In the Contents area, enter the following commands:

```
#Title: red action file for Partition Utilization agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example, if the agent is monitoring all classes you can send an email notification that indicates the number of classes whose partitions had utilization higher than the Utilization threshold during the evaluation interval:

```
send email raltman@test.com "High Partition Utilization" "$scorevalue classes had
utilization greater than $Utilization percent."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

If the agent is monitoring a single class, you can include the variables for the class name ($class-id) and the utilization ($avg-bps) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "$class-id partition" "Average utilization for the evaluation
interval was $avg-bps bps. The utilization for $class-id exceeded $Utilization percent of
the partition size."
```

When the message is sent, the body of the email message would look something like this:

```
Average utilization for the evaluation interval was 383047 bps. The utilization for
/Inbound/ftp exceeded 90 percent of the partition size.
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *ftp-part.cmd*).

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed. If you like, you can create a green action file to notify you when the Partition Utilization agent returns to a green status.

Analyze Efficiency with Traffic Performance Agents

Network efficiency is the percentage of bytes that were not retransmissions. A network efficiency of 100 percent means that no bytes were retransmitted in other words, the network was highly efficient. The Network Performance Summary on the Report tab is one way to check the efficiency of your link. Another way is to use a Traffic Performance agent. You can monitor the network efficiency of your Inbound and Outbound links, a specific traffic class, or all classes.

When a class' efficiency dips below a certain percentage (such as 80 percent), the adaptive response feature will alert you that a threshold has been crossed. The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when the class' efficiency drops under a predefined threshold; this method requires that you create an *action file*– a text file of PacketShaper commands. The action file can take other actions, as well, such as enabling top talkers/listeners. The action file can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to use this feature, the adaptive response feature must be enabled and you must create an agent based on the Traffic Performance template.

Create a Traffic Performance Agent

To create a Traffic Performance agent that monitors the efficiency of one or more classes:

1. On the **Setup** tab, select the **Adaptive Response** setup page.
2. Click **add**. The *Add Agent Entry* window opens.
3. In the **create a new agent from this agent template** drop-down list, choose *Traffic Performance*.
4. Enter a unique identifying name for the new agent in the **Name** field, for example *Inbound Efficiency* or *Oracle Efficiency*.
5. If desired, change the evaluation interval, in minutes (how often the agent should measure its target).
6. Click **OK and Edit Parms**. This will open the *Edit Agent Entry* window where you can define the parameter values for your new agent.
7. In the **ClassName** field, enter the name of the class to be monitored (for example, */Inbound/Oracle*). If desired, a wildcard may be used to monitor many classes (for example, */inbound** to monitor all classes in the inbound direction or * to monitor all classes).

> **Note:** This class must already exist in your traffic tree.

8. If desired, modify the threshold values:

   Efficiency– The percent efficiency that will cause a class to be in violation. For instance, if the Efficiency is set to 80 and the efficiency of one of the monitored classes drops below 80 percent, that class will be in violation. Assuming the RedThreshold is set to 1 (the default), the agent's status will change to red.

   RedThreshold– The value at which the status of the agent turns red, and a red action file is triggered, if present. For a Traffic Performance agent, this number represents an unacceptable number of violating classes; a class is in violation if it is less than the defined Efficiency percentage. For instance, if the RedThreshold is set to 1, the agent's status will change to red when at least one class had efficiency less than the defined Efficiency percent.

   GreenThreshold– An acceptable number of violating classes (usually 0).

9. Click **OK** to save your changes.

After creating the agent, you will want to monitor it to see whether the class' efficiency ever dropped below the acceptable level you defined.

Check the Agent Status

To check the status of Traffic Performance agents:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. Traffic Performance agents are in the Network Health category (the first one).

2. To see the status of each individual agent in the Network Health category, hover your mouse over the first colored status indicator, as shown below.

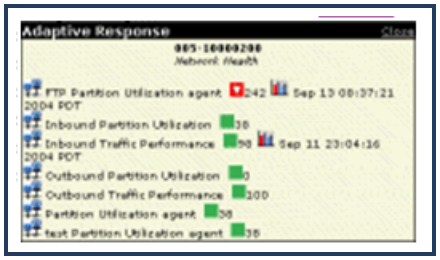

3. In the pop-up window, locate the traffic performance agent that you defined. What color is its status indicator?

   Green– During the last evaluation interval, efficiency of the defined class name(s) did not drop below the Efficiency percentage you defined. In other words, traffic was efficient few packets were dropped.
   Red– During the last evaluation interval, efficiency dropped below the defined Efficiency percentage. In other words,

retransmitted packets caused inefficiency in the class(es). In this case, you may want to consider delving into this further, perhaps checking the class' response times. When the status indicator is red, you should look at the incident report for details.

If the agent has ever changed status (from green/yellow to red or from red/yellow to green, a report icon with a link to a detailed report will be available for that agent.

4. If there is a report for the Traffic Performance agent, click the report icon. The incident report appears in a separate browser window and lists all the classes during the interval that had efficiency less than the specified Efficiency percentage.

Incident Report for Traffic Performance agent at Thu Jan 13 05:42:01 2005 PST

"time","agent4","network-efficiency","class-id"
"13-Jan-2005 05:41:00","Traffic Performance agent",80,/Inbound/Oracle
"13-Jan-2005 05:42:00","Traffic Performance agent",77,/Inbound

Unit IP Address: 172.21.18.160 Name: myshaper

Create an Action File that Sends an Email Notification

If you would like to be notified by email when classes are inefficient, you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file that sends an email notification:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and select the **Adaptive Response** setup page.

2. Click the **edit** button next to the agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, A-Z, -,_, and . (period). Spaces are not allowed. Example: *rtx-red*

6. In the Contents area, enter the following commands:

```
#Title: red action file for Traffic Performance Agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example, if the agent is monitoring all classes you can send an email notification that indicates the number of classes that were inefficient during the evaluation interval:

```
send email raltman@test.com "Low efficiency" "$scorevalue classes had efficency less than
$Efficiency percent."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

If the agent is monitoring a single class, you can include the variables for the class name ($class-id) and the efficiency ($network-efficiency) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "Low $class-id efficiency" "The efficiency for $class-id was
$network-efficiency percent."
```

When the message is sent, the body of the email message would look something like this:

```
The efficiency for inbound/oracle was 79 percent.
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *rtx-red.cmd).*

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed. If you like, you can create a green action file to notify you when the agent returns to a green status.

Use the Memory Allocation Agent

The Memory Allocation agent monitors memory failures on your PacketShaper. The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when the system load is high (such as 95%); this method requires that you create an *action file*– a text file of PacketShaper commands. The action file can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to use this feature, adaptive response must be enabled and you must create an agent based on the Memory Allocation template.

Create a Memory Allocation Agent

To create a Memory Allocation agent:

1. On the **Setup** tab, click the **Adaptive Response Setup** link.
2. Click **add**. The *Add Agent Entry* window will open.
3. In the **create a new agent from this agent template** drop-down list, choose *Memory Allocation*.
4. Enter a unique identifying name for the new agent in the **Name** field.
5. If desired, change the evaluation interval, in minutes (how often the agent should measure its target).
6. Click **OK and Edit Parms**. This will open the *Edit Agent Entry* window where you can define the parameter values for your new agent.
7. If desired, modify the threshold values:

   RedThreshold The value at which the status of the agent turns red, and a red action file is triggered, if present. For instance, if the RedThreshold is set to 2, the agent's status will change to red when there are more than 2 memory failures in the interval.

   GreenThreshold The value at which the status of the agent turns green, and a green action file is triggered, if present. For instance, if the GreenThreshold is set to 0, the agent's status will be green when the unit doesn't have any memory allocation problems in the interval.

8. Click **OK** to save your changes.

Check the Agent Status

> **Note:** Before you can check the status of the Memory Allocation agent, you will need to wait until the first evaluation interval has passed (60 minutes, by default).

To check the status of the Memory Allocation agent:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The Memory Allocation agent is in the Unit Health category.

2. To see the status of the Memory Allocation agent, hover your mouse over the Unit Health colored status indicator, as shown below.

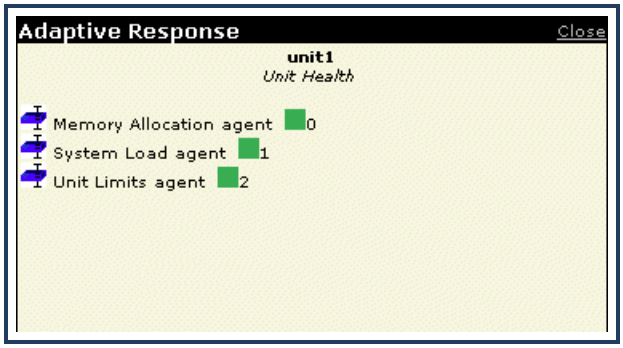

3. In the pop-up window, locate the Memory Allocation agent. The value (0 in the figure above) indicates the number of memory failures in the interval.

   What color is its status indicator?

   Green– During the last evaluation interval, the PacketShaper did not have any memory allocation errors.
   Yellow– During the last evaluation interval, the unit had at least one memory failure. (Yellow is the range between the red and green thresholds.
   Red– During the last evaluation interval, the PacketShaper experienced several memory failures (more than the defined red threshold). If the status indicator is red, you should look at the incident report for details.

   If the agent has ever changed status (from green/yellow to red or from red/yellow to green, a report icon with a link to a detailed incident report will be available for that agent.

4. If there is a report for the Memory Allocation agent, click the report icon.

Create an Action File

If you would like to be notified by email when the PacketShaper is experiencing memory failures above the red threshold, you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and click the **Adaptive Response Setup** link.

2. Click the **edit** button next to the Memory Allocation agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, A-Z, -,_. Spaces are not allowed. The file will automatically be assigned the extension .cmd. Example: *red-mem*

6. In the Contents area, enter the following commands:

```
#Title: red action file for Memory Allocation agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example:

```
send email raltman@test.com "Memory Allocation" "The PacketShaper is currently
experiencing problems with memory allocation. Go to the Info tab to view the incident
report for the Memory Allocation agent."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

Alternatively, you can include the variables for the agent name ($agentname) and the number of memory failures ($scorevalue) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "$agentname agent" "The PacketShaper recently had $scorevalue
memory failures. Go to the Info tab to view the incident report for the $agentname agent."
```

When the message is sent, the body of the email message would look something like this:

```
The PacketShaper recently had 3 memory failures. Go to the Info tab to view the incident
report for the Memory Allocation agent.
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *red-mem.cmd)*.

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed. If you like, you can create a green action file to notify you when the Memory Allocation agent returns to a green status.

Use the System Load Agent

The System Load agent closely monitors the system processing load on your PacketShaper and alerts you when the load becomes too high. The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when the system load is high (such as 95%); this method requires that you create an *action file*– a text file of PacketShaper commands. The action file can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to use this feature, adaptive response must be enabled and you must create an agent based on the System Load template.

Create a System Load Agent

To create a System Load agent:

1. On the **Setup** tab, choose the **Adaptive Response** setup page.
2. Click **add**. The *Add Agent Entry* window opens.
3. In the **create a new agent from this agent template** drop-down list, choose *System Load*.
4. Enter a unique identifying name for the new agent in the **Name** field.
5. If desired, change the evaluation interval, in minutes (how often the agent should measure its target).
6. Click **OK and Edit Parms**. This will open the *Edit Agent Entry* window where you can define the parameter values for your new agent.
7. If desired, modify the threshold values:

   RedThreshold– The value at which the status of the agent turns red, and a red action file is triggered, if present. For the System Load agent, this number is the percentage at which you consider the unit's CPU utilization to be too high. For instance, if the RedThreshold is set to 95, the agent's status will change to red when the CPU utilization is greater than 95 percent.

   GreenThreshold– The value at which the status of the agent turns green, and a green action file is triggered, if present. For instance, if the GreenThreshold is set to 90, the agent's status will be green when the unit's system load is less than 90 percent.
8. Click **OK** to save your changes.

Check the Agent Status

To check the status of the System Load agent:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The System Load agent is in the Unit Health category.

2. To see the status of the System Load agent, hover your mouse over the Unit Health colored status indicator, as shown below.

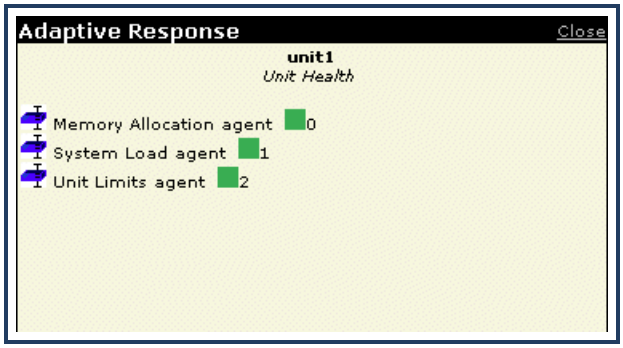

3. In the pop-up window, locate the System Load agent. The value indicates the unit's CPU utilization. For example, a score of 67 indicates the unit's CPU utilization was 67 percent (33 percent idle) in the last evaluation interval.

   What color is its status indicator?

   Green– During the last evaluation interval, the PacketShaper was not being overutilized.
   Yellow– During the last evaluation interval, the unit's CPU utilization was getting close to the red threshold. (Yellow is the range between the red and green thresholds.)
   Red– During the last evaluation interval, the PacketShaper's CPU utilization was getting too high. If the status indicator is red, you should look at the incident report for details.

   If the agent has ever changed status (from green/yellow to red or from red/yellow to green, a report icon with a link to a detailed incident report will be available for that agent.

4. If there is a report for the System Load agent, click the report icon. The incident report looks similar to the one below.

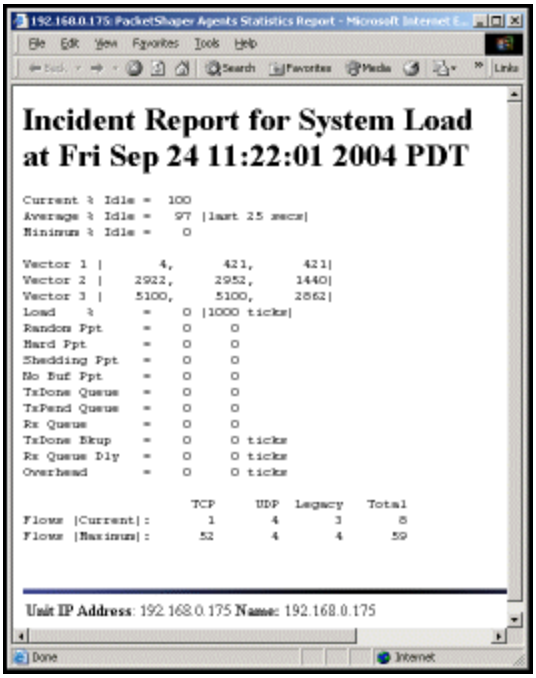

The System Load agent's incident report provides information about CPU utilization the % idle is the amount the CPU is *not* being utilized. For example, if the Average % Idle is 5, the CPU utilization is 95%. The middle section contains a list of internal buffers; high numbers are indicative that the PacketShaper is under a heavy load, and are at an inappropriate level for the capacity of the unit. The bottom section lists information about the flows going through the unit. If these values are close to the unit's configuration limits, the unit can get overloaded.

Create an Action File

If you would like to be notified by email when the PacketShaper's system load is high (above the red threshold), you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and choose the **Adaptive Response** setup page.

2. Click the **edit** button next to the System Load agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, A-Z, -,_. Spaces are not allowed. The file will automatically be assigned the extension .cmd. Example: *red-load.cmd*

6. In the Contents area, enter the following commands:

```
#Title: red action file for System Load agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example:

```
send email raltman@test.com "System Load" "The PacketShaper's system load is currently very high. Go to the Info tab to view the incident report for the System Load agent."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

Alternatively, you can include the variables for the agent name ($agentname) and the CPU utilization ($scorevalue) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "$agentname agent" "The PacketShaper's system load is currently very high ($scorevalue percent). Go to the Info tab to view the incident report for the $agentname agent."
```

When the message is sent, the body of the email message would look something like this:

```
The PacketShaper's system load is currently very high (96 percent). Go to the Info tab to view the incident report for the System Load agent.
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *red-load.cmd*).

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed. If you like, you can create a green action file to notify you when the System Load agent returns to a green status.

Use the Unit Limits Agent to Track Capacity

The Unit Limits agent carefully monitors the number of classes, matching rules, partitions, and dynamic partitions on your PacketShaper and alerts you when the unit is approaching its limits. Note that the agent does not report specific values for each of these objects it tracks a composite score.

The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when the unit is close to system capacity (such as 99 percent); this method requires that you create an *action file*– a text file of PacketShaper commands. The action file can take other actions, as well, such as disabling traffic discovery or deleting unnecessary classes (for example, port numbered classes). The file can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to use this feature, adaptive response must be enabled and you must create an agent based on the Unit Limits template.

Create a Unit Limits Agent

By default, the Unit Limits agent checks the current number of classes, matching rules, partitions, and dynamic partitions every 15 minutes. It compiles this information into a composite score, represented as a percentage of system capacities. With the default settings, a score greater than 99 percent would trigger a red alert.

To create a Unit Limits agent:

1. On the **Setup** tab, choose the **Adaptive Response** setup apge.
2. Click **add**. The *Add Agent Entry* window opens.
3. In the **create a new agent from this agent template** drop-down list, choose *Unit Limits*.
4. Enter a unique identifying name for the new agent in the **Name** field.
5. If desired, change the evaluation interval, in minutes (how often the agent should measure its target).
6. Click **OK and Edit Parms**. This will open the *Edit Agent Entry* window where you can define the parameter values for your new agent.
7. If desired, modify the threshold values:

   RedThreshold –The value at which the status of the agent turns red, and a red action file is triggered, if present. For the Unit Limits agent, this number represents the percentage at which the unit is getting close to maximum system capacities. For instance, if the RedThreshold is set to 95, the agent's status will change to red when the composite score is greater than 95 percent of system capacity.

GreenThreshold– The value at which the status of the agent turns green, and a green action file is triggered, if present. For instance, if the GreenThreshold is set to 90, the agent's status will be green when the unit's composite score is less than 90 percent of system capacity.

8. Click **OK** to save your changes.

Check the Agent Status

> **Note:** Before you can check the status of the Unit Limits agent, you will need to wait until the first evaluation interval has passed (15 minutes, by default).

To check the status of the Unit Limits agent:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The Unit Limits agent is in the Unit Health category.

2. To see the status of the Unit Limits agent, hover your mouse over the Unit Health colored status indicator, as shown below.

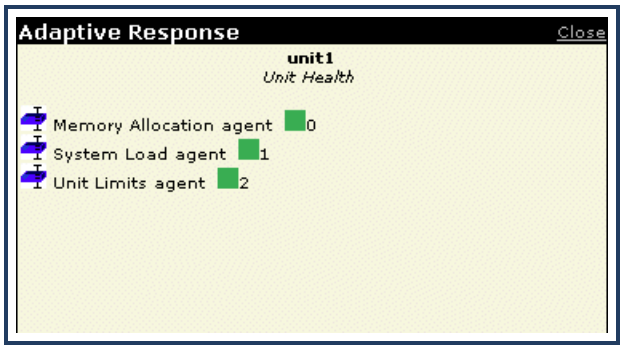

3. In the pop-up window, locate the Unit Limits agent. The value (6 in the figure above) indicates the unit's composite score represented as a percentage of system capacities. For example, a score of 63 indicates the unit was at 63 percent of capacity in the last evaluation interval.

   What color is its status indicator?

   Green– During the last evaluation interval, the PacketShaper was not close to reaching maximum capacity.
   Yellow– During the last evaluation interval, the composite score was getting close to the red threshold. (Yellow is the range between the red and green thresholds.)
   Red– During the last evaluation interval, the PacketShaper was approaching maximum capacity. If the status indicator is red, you should look at the incident report for details.

If the agent has ever changed status (from green/yellow to red or from red/yellow to green, a report icon with a link to a detailed incident report will be available for that agent.

4. If there is a report for the Unit Limits agent, click the report icon. The incident report lists the maximum number of objects allowed, currently used, and remaining. You can use this report to see which objects are approaching their limits. For example, if the number of classes is approaching capacity, you might want to delete unnecessary classes from the traffic tree.

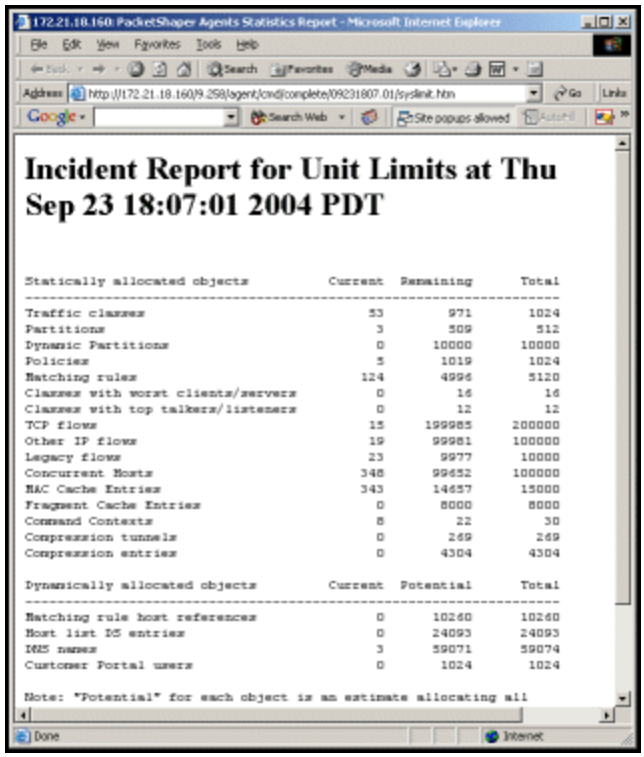

Create an Action File

If you would like to be notified by email when the PacketShaper is approaching maximum capacity, you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and choose the **Adaptive Response** setup page.

2. Click the **edit** button next to the Unit Limits agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, AZ, -,_. Spaces are not allowed. The file will automatically be assigned the extension .cmd. Example: *red-lmt*

6. In the Contents area, enter the following commands:

```
#Title: red action file for Unit Limits agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example:

```
send email raltman@test.com "Unit Limits" "The PacketShaper is currently close to maximum
capacity. Go to the Info tab to view the incident report for the Unit Limits agent."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

Alternatively, you can include the variables for the agent name ($agentname) and the composite score ($scorevalue) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "$agentname agent" "The PacketShaper is at $scorevalue of its
capacity. Go to the Info tab to view the incident report for the $agentname agent."
```

When the message is sent, the body of the email message would look something like this:

```
The PacketShaper is at 96 percent of its capacity. Go to the Info tab to view the incident
report for the Unit Limits agent.
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *red-lmt.cmd*).

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed. After reviewing the incident report, you can decide on the appropriate course of action, such as deleting unnecessary classes, turning off auto-discovery, expanding the capacity of your unit by purchasing a software key, or upgrading to a higher-capacity unit. If you are running out of matching rules, it's possible that the rules have not been defined as efficiently as they could be. For instance, instead of creating separate matching rules for each host, you can create a single matching rule with a host list.

If you like, you can create a green action file to notify you when the Unit Limits agent returns to a green status.

Create Default Traffic Agents

The Default Traffic agent alerts you when the traffic in a default traffic class (such as Inbound/Default) exceeds a certain percentage of the parent class partition in a specific time interval. For example, if more than 15 percent of the parent class partition is not getting classified and therefore being placed into the /default class, you would want to know about this situation and perhaps take action on it. What could be causing a Default class to have so much traffic? You may have turned off traffic discovery, perhaps because you have fine tuned your traffic tree with just the classes you care about monitoring and controlling; if new applications are being run on your network, they will get classified into Default. Another possible reason for a large amount of traffic in a Default class is that you have auto-discovery on, but the PacketShaper has created the maximum number of classes no more classes can be discovered so the traffic gets classified into Inbound/Default or Outbound/Default.

The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when the Default class' utilization crosses a predefined threshold; this method requires that you create an *action file* – a text file of PacketShaper commands. The action file can take other actions, as well, such as enabling traffic discovery or turning on top talkers/listeners, and can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to create agents, the adaptive response feature must be enabled.

Create a New Agent

To create a Default Traffic agent:

> **Note:** If you want to monitor Inbound/Default or Outbound/Default, you don't need to create a new agent. You can simply check the status of the predefined Inbound Default Traffic and Outbound Default Traffic agents.

1. On the **Setup** tab, choose the **Adaptive Response** setup page.
2. Click **add**. The *Add Agent Entry* window opens.
3. Click the **create a new agent from this agent template** drop-down list, and choose *Default Traffic*.
4. Enter a unique identifying name for the new agent in the **Name** field, for example *Citrix Default Traffic*.
5. If desired, change the evaluation interval, in minutes (how often the agent should measure its target).
6. Click **OK and Edit Parms**. This will open the *Edit Agent Entry* window where you can define the parameter values for your new agent.

7. In the **ClassName** field, enter the name of the class to be monitored (for example, *Inbound/Citrix/Default*). Note that this must be a Default class.

8. If desired, modify the threshold values:

   RedThreshold– The value at which the status of the agent turns red, and a red action file is triggered, if present. For a Default Traffic agent, this number represents the Default class' percent utilization of the parent class partition during the evaluation interval. For instance, if the RedThreshold is set to 15, the agent's status will change to red when the utilization of the class is more than 15 percent of the parent class partition.

   GreenThreshold – The value at which the status of the agent turns green, and a green action file is triggered, if present. For instance, if the GreenThreshold is set to 7, the agent's status will be green when the utilization of the Default class is less than 7 percent of the parent class partition.

9. Click **OK** to save your changes.

After creating the agent, you will want to monitor it to see whether usage in the class has come close to or exceeded the threshold.

Check the Agent Status

To check the status of the Default Traffic agent:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The Default Traffic agents are in the Application Health category.

2. To see the status of each individual agent in the Application Health category, hover your mouse over the colored status indicator, as shown below.

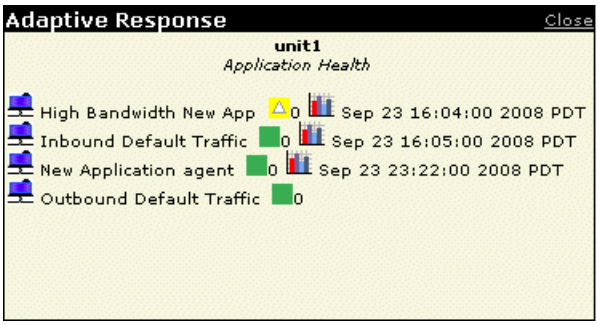

3. In the pop-up window, locate the Default Traffic agents. What colors are its status indicators?

   Green– During the last evaluation interval, usage did not exceed the red threshold you defined. In other words, an acceptable amount of traffic is being classified into the Default class.
   Yellow– During the last evaluation interval, usage was approaching the red threshold (yellow is the range between the red

and green thresholds).
Red – During the last evaluation interval, usage exceeded the defined red threshold. In other words, too much traffic is getting classified into Default. In this situation, you may want to consider taking some action, such as enabling traffic discovery or turning on top talkers/listeners. If the status indicator is red, you should look at the incident report for details.

If the agent has ever changed status (from green/yellow to red or from red/yellow to green, a report icon with a link to a detailed report will be available for that agent.

4. If there is a report for the Default Traffic agent, click the report icon. The incident report appears in a separate browser window.

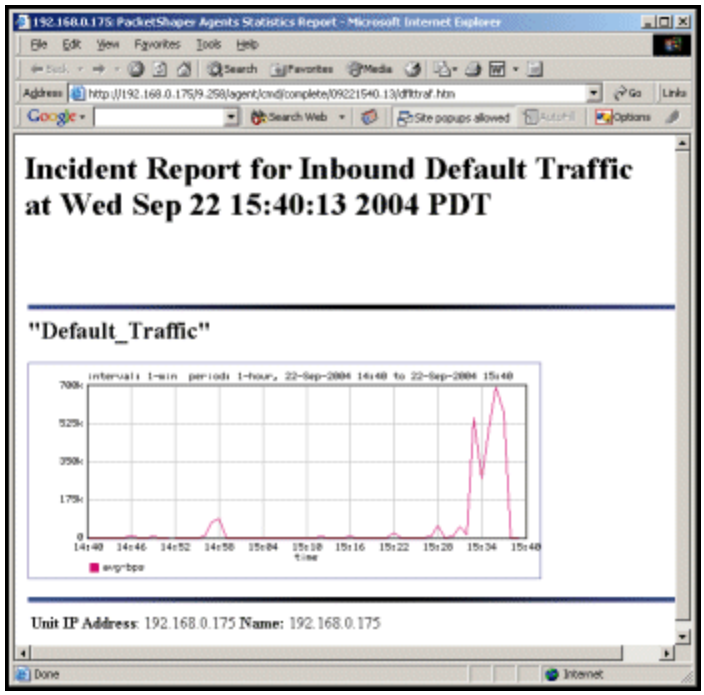

Create an Action File that Sends an Email Notification

If you would like to be notified by email when the Default class' usage has exceeded the red threshold, you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed. In addition, the action file can contain CLI commands for taking corrective action (enabling traffic discovery) or diagnostic action (turning on top talkers/listeners).

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file that sends an email notification:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and click the **Adaptive Response Setup** link.

2. Click the **edit** button next to the partition agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, AZ, -,_, and . (period). Spaces are not allowed. Example: *red-def*

6. In the Contents area, enter the following commands:

```
#Title: red action file for Default Traffic agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example:

```
send email raltman@test.com "Inbound/Default class exceeds threshold" "The utilization for
inbound/default exceeded 15% of the parent class partition."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

Alternatively, you can include the variables for the class name ($class-id), average utilization ($avg-bps), and the red threshold ($RedThreshold) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "$class-id exceeds threshold" "Average utilization for the
evaluation interval was $avg-bps bps. The utilization for $class-id exceeded $RedThreshold
percent of the parent class partition."
```

When the message is sent, the body of the email message would look something like this:

```
Average utilization for the evaluation interval was 383047 bps. The utilization for
/Inbound/default exceeded 15 percent of the parent class partition.
```

If you want the action file to enable traffic discovery, include the following command in the file:

```
setup discover on
```

Or, if you want the action file to turn on top talkers/listeners for the Default class:

```
hostdb topusers start $class-id listen
 hostdb topusers start $class-id talk
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *red-def.cmd*).

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed.

If the action file turned on top talkers/listeners for the class, you will want to wait a period of time for PacketShaper to gather host information. Then you can view the Top Talkers and Top Listeners reports for the class to discover the hosts that are the top bandwidth consumers; from this information you might decide to manually create classes for some of these hosts, in order to reduce the amount of unclassified traffic.

You may also want to create a green action file to notify you when the Default Traffic agent returns to a green status. This action file could turn off traffic discovery (**setup discover off**) or turn off top talkers/listeners (**hostdb topusers stop $class-id talk**).

Use the New Application Agent

You can use the New Application agent to learn when an excessive number of classes are being created in your traffic tree with the auto-discovery feature. This agent is useful after the initial creation of your tree (you wouldn't want to enable the agent while the tree is in the process of being built). This agent can alert you of a possible threat to your network. For example, if someone was using a port scanner to look for open ports on your network to gain illegal entry, a large number of port numbered classes would be created in a short period of time.

The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when the number of auto-discovered classes crosses a predefined threshold; this method requires that you create an *action file* – a text file of PacketShaper commands. The action file can take other actions, as well, such as disabling traffic discovery, and can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to use this feature, the adaptive response feature must be enabled and you must create an agent based on the New Application template.

Note that this agent will always return a "green" status when traffic discovery is enabled. The New Application agent is intended to be used in conjunction with traffic discovery. If discovery is always disabled on your unit, you can disable this agent.

Create an Agent Based on the New Application Template

The New Application agent triggers "red" when the number of new auto-discovered classes exceeds the red threshold. However, the default red threshold is so high (99999 classes) that it will always report a green status, regardless of the number of new applications discovered in the interval. Therefore, you will need to modify the default settings to define your own red threshold according to what you think would be an excessive number of new classes to be discovered on your network during a set period of time. The default evaluation interval is one day (1440 minutes). You can adjust the duration if you want the agent to calculate the number of newly discovered applications more frequently.

To create a New Application agent:

1. On the **Setup** tab, choose the **Adaptive Response** setup page.
2. Click **add**. The *Add Agent Entry* window opens.
3. In the **create a new agent from this agent template** drop-down list, choose *New Application*.
4. Enter a unique identifying name for the new agent in the **Name** field.
5. If desired, change the evaluation interval, in minutes (how often the agent should measure its target).

6. Click **OK and Edit Parms**. This will open the *Edit Agent Entry* window where you can define the parameter values for your new agent.

7. If desired, modify the threshold values:

   RedThreshold– The value at which the status of the agent turns red, and a red action file is triggered, if present. For the New Application agent, this number represents the number of new auto-discovered classes considered to be excessive. For instance, if the RedThreshold is set to 50, the agent's status will change to red when there are more than 50 new auto-discovered classes in the interval.

   GreenThreshold – The value at which the status of the agent turns green, and a green action file is triggered, if present. For instance, if the GreenThreshold is set to 40, the agent's status will be green when fewer than 40 classes are auto-discovered in the interval.

8. Click **OK** to save your changes.

After creating the agent, you will want to monitor it to see whether the number of new auto-discovered classes has come close to or exceeded the threshold.

Check the Agent Status

> **Note:** Before you can check the status of the New Application agent, you will need to wait until the first evaluation interval has passed (1440 minutes, by default).

To check the status of the New Application agent:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The New Application agent is in the Application Health category.

2. To see the status of each individual agent in the Application Health category, hover your mouse over the colored status indicator, as shown below.

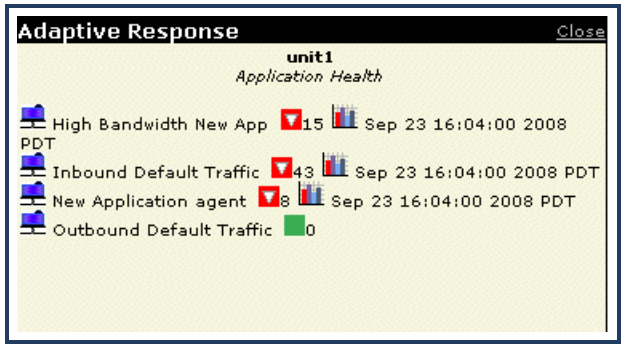

3. In the pop-up window, locate the New Application agent. The value (20 in the figure above) indicates the number of auto-discovered classes during the interval.

   What color is its status indicator?

   Green– During the last evaluation interval, the number of new auto-discovered classes did not exceed the red threshold you defined. In other words, the number of classes being auto-discovered is not alarming.
   Yellow– During the last evaluation interval, the number of new auto-discovered classes was approaching the red threshold. (Yellow is the range between the red and green thresholds.)
   Red– During the last evaluation interval, the number of new auto-discovered classes exceeded the defined red threshold. In other words, an excessive number of classes were discovered and created during the interval.

   If the agent has ever changed status (from green/yellow to red or from red/yellow to green), a report icon with a link to a detailed report will be available for the agent.

4. If there is a report for the New Application agent, click the report icon. The incident report lists all the classes auto-discovered during the evaluation interval.

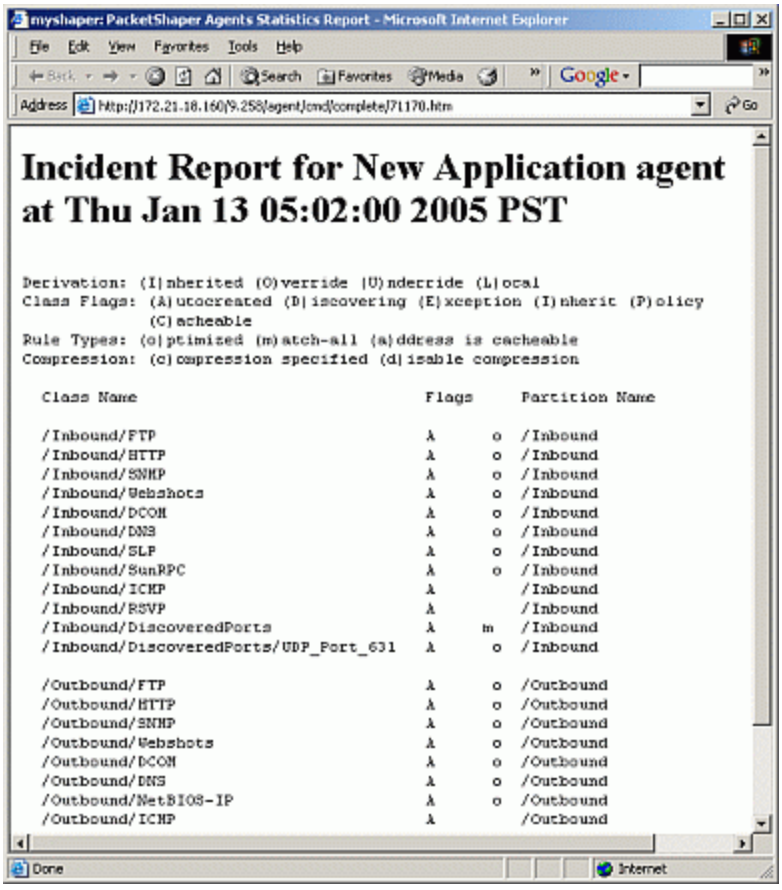

Create an Action File that Sends an Email Notification

If you would like to be notified via email when the number of new auto-discovered classes has exceeded the red threshold, you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file that sends an email notification:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and choose the **Adaptive Response** setup page.

2. Click the **edit** button next to the partition agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, AZ, -,_, and . (period). Spaces are not allowed. Example: *red-new*

6. In the Contents area, enter the following commands:

```
#Title: red action file for New Application agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example:

```
send email raltman@test.com "New discovered classes" "$scorevalue applications were auto-discovered during the last interval."
```

The $scorevalue variable represents the number of classes that were auto-discovered during the interval. When the message is sent, the body of the email message would look something like this:

```
 20 applications were auto-discovered during the last interval.
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *red-new.cmd*).

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed.

You may also want to create a green action file to notify you when the New Application agent returns to a green status.

Monitor Bandwidth of New Applications

The High Bandwidth New App agent monitors the bandwidth usage of newly discovered classes on your network. If the traffic in a single new class exceeds a certain percent of the parent partition size, the adaptive response feature will alert you that a threshold has been exceeded. The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when a high bandwidth application has been found; this method requires that you create an *action file* – a text file of PacketShaper commands. Taking it a few steps further, you can have the action file automatically create a partition to cap the bandwidth usage in the class. The action file can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to use this feature, adaptive response and traffic discovery must be enabled and the High Bandwidth New App agent must be turned on.

Modify the Agent Settings

By default, the High Bandwidth New App agent checks hourly to see if a single class is using more than 10 percent of the parent partition's bandwidth (typically the Inbound or Outbound link size). If you want to change the evaluation interval (for instance, to every 30 minutes) or change the percentage at which an application is considered to be using too much bandwidth, you can edit the agent's settings.

1. On the **Setup** tab, choose the **Adaptive Response** setup page.

2. Make sure the check box in the Agent On column is selected for the High Bandwidth New App agent.

3. Click **edit** next to the High Bandwidth New App agent. The *Edit Agent Entry* window will open.

4. If desired, change the evaluation interval, in minutes (how often the agent should measure its target.)

5. If desired, modify the threshold values:

   RedThreshold The value at which the status of the agent turns red, and a red action file is triggered, if present. For the High Bandwidth New App agent, this number represents the percentage at which a newly discovered class is considered to be using too much bandwidth. For instance, if the RedThreshold is set to 10, the agent's status will change to red when the utilization of a single class is more than 10 percent of the parent partition size.

   GreenThreshold The value at which the status of the agent turns green, and a green action file is triggered, if present. For instance, if the GreenThreshold is set to 5, the agent's status will be green when no single class has more than five percent utilization of its parent partition.

6. Click **OK** to save your changes.

Check the Agent Status

To check the status of the High Bandwidth New App agent:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The High Bandwidth New App agent is in the Application Health category.

2. To see the status of the High Bandwidth New App agent, hover your mouse over the Application Health colored status indicator, as shown below.

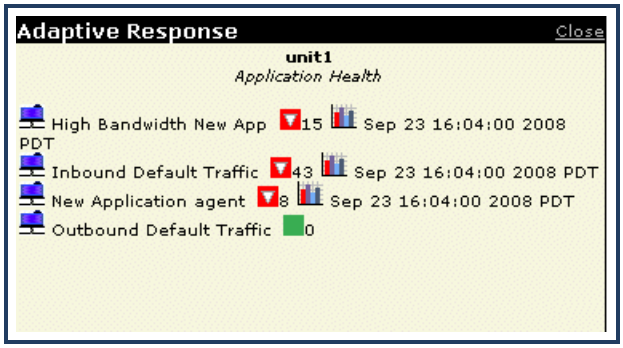

3. In the pop-up window, observe the color of the High Bandwidth New App agent's status indicator.

   Green During the last evaluation interval, usage of any single newly discovered class did not exceed the red threshold you defined. That is, no new class used a disproportionate share of bandwidth during the interval.
   Yellow During the last evaluation interval, usage was approaching the red threshold (yellow is the range between the red and green thresholds).
   Red During the last evaluation interval, usage exceeded the defined red threshold. In other words, at least one new class was using too much bandwidth. If the status indicator is red, you should look at the incident report for details.

   If the agent has ever changed status (from green/yellow to red or from red/yellow to green, a report icon with a link to a detailed report will be available for that agent.

4. If there is a report for the High Bandwidth New App agent, click the report icon. The incident report lists the name of the class that was consuming excessive bandwidth, the average bandwidth used in the interval, and a utilization graph of this class.

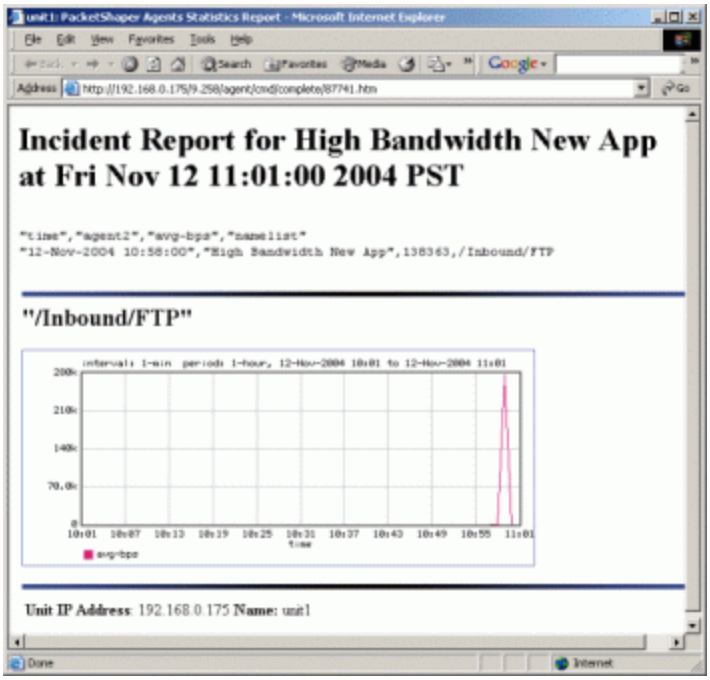

Create an Action File

There are different actions you might want to take if a High Bandwidth New App is found on your network. Perhaps you just want an email notification that indicates the name of the application that is consuming excessive bandwidth. Or perhaps you want to create a partition to limit the bandwidth used by the class.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and click the **Adaptive Response Setup** link.

2. Click the **edit** button next to the High Bandwidth New App agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, A-Z, -,_, and . (period). Spaces are not allowed. Example: *red-app*

6. In the Contents area, enter the following commands:

```
#Title: red action file for High Bandwidth New App agent
 partition apply $namelist 5% fixed
```

```
 send email <address> "High Bandwidth Consuming New App found" "New class $namelist
recently used $avg-bps bps, which is more than the Red Threshold ($RedThreshold percent). A
partition was applied to Class $namelist."
```

where *<address>* is the email address of the recipient. The $namelist variable represents the name of the class that is using the most bandwidth in the interval. The $avg-bps variable is the class' utilization, in bits per second. To summarize, this action file will apply a partition to the class, restricting the amount of bandwidth available to the class; in this example, the partition size is 5% of the parent partition, but you can set any appropriate size. You will then be sent an email notifying that these actions have been taken.

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

When the message is sent, the body of the email message would look something like this:

```
New class inbound/ftp recently used 762744 bps, which is more than the Red Threshold (10
percent). A partition was applied to Class $namelist.
```

If you simply want to receive an email notification that a new class is using excessive bandwidth, the action file would look like this:

```
#Title: red action file for High Bandwidth New App agent
 send email <address> "High Bandwidth Consuming New App found" "New class $namelist
recently used $avg-bps bps, which is more than the Red Threshold ($RedThreshold percent)."
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *red-app.cmd*).

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed.

Identify and Control High Bandwidth Hosts

Adaptive response offers two agents the track the bandwidth usage of individual hosts on the network. With the High Bandwidth Host agent, the adaptive response feature will alert you if a single host's traffic exceeds a certain percent of the link size. With the Quota Bandwidth Host agent, adaptive response will alert you if a single host's absolute bandwidth consumption exceeds a defined value for a specified interval.

The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when a bandwidth abusing host has been found; this method requires that you create an *action file* – a text file of PacketShaper commands. Taking it a few steps further, you can have the action file automatically create a traffic class for the offending host so that you can track and monitor it more closely. The action file can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to use this feature, adaptive response must be enabled and you must create an agent based on the High Bandwidth Host or Quota Bandwidth Host template.

Create a High Bandwidth Host Agent

To create a High Bandwidth Host agent:

1. On the **Setup** tab, choose the **Adaptive Response** setup page.
2. Click **add**. The *Add Agent Entry* window opens.
3. In the **create a new agent from this agent template** drop-down list, choose *High Bandwidth Host*.
4. Enter a unique identifying name for the new agent in the **Name** field.
5. If desired, change the evaluation interval, in minutes (how often the agent should measure its target).
6. Click **OK and Edit Parms**. This will open the *Edit Agent Entry* window where you can define the parameter values for your new agent.
7. If desired, modify the threshold values:

   RedThreshold– The value at which the status of the agent turns red, and a red action file is triggered, if present. For the High Bandwidth Host agent, this number represents the percentage at which a single host is considered to be using too much bandwidth. For instance, if the RedThreshold is set to 10, the agent's status will change to red when the utilization of a single host is more than 10 percent of the link size.

   GreenThreshold – The value at which the status of the agent turns green, and a green action file is triggered, if

present. For instance, if the GreenThreshold is set to 5, the agent's status will be green when no single host has more than five percent utilization of the link.

8. Click **OK** to save your changes.

Create a Quota Bandwidth Host Agent

To create a Quota Bandwidth Host agent:

1. On the **Setup** tab, click the **Adaptive Response Setup** link.

2. Click **add**. The *Add Agent Entry* window will open.

3. In the **create a new agent from this agent template** drop-down list, choose *Quota Bandwidth Host*.

4. Enter a unique identifying name for the new agent in the **Name** field.

5. If desired, change the evaluation interval, in minutes (how often the agent should measure its target).

6. Click **OK and Edit Parms**. This will open the *Edit AgentEntry* window where you can define the parameter values for your new agent.

7. If desired, modify the agent parameters:

   HostUsageThreshold – Maximum number of bytes a host may use during the interval. Threshold can be between 0 and 2,000,000,000 bytes.

   HostUsageMonitorInterval – Interval to be monitored (in days). Interval can be between 0 and 30 days. At the end of the interval, the measured bytes for each host is zeroed out and any violating hosts are removed from the ViolatingHosts list. Note: The interval begins when the agent is enabled. To control the time of day at which the bytes are reset (such as at midnight), you'll need to enable the agent at the desired time.

   Side – The traffic direction to be monitored: inside, outside, or both

   ViolatingHosts – Name of host list. Adaptive response will automatically add violating hosts to this host list. For more information about using violating host lists, see Violating Host Lists.

   ExceptionHosts – Name of host list that contains hosts to be excluded. You should edit this host list and add the IP addresses of any hosts (such as servers) that you don't need to monitor. See Host Exception Lists.

   RedThreshold – Leave this at the default value (1). If one or more hosts exceed the HostUsageThreshold during the evalutation interval, the agent will be in violation (red).

GreenThreshold – Leave this parameter at the default value (0). The agent will be green when no hosts exceed the HostUsageThreshold during the evaluation interval.

8. Click **OK** to save your changes.

Check the Agent Status

To check the status of your Bandwidth Host agent:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The High Bandwidth Host and Quota Bandwidth Host agents are in the Hosts category.

2. To see the status of the Bandwidth Host agent, hover your mouse over the Hosts colored status indicator, as shown below.

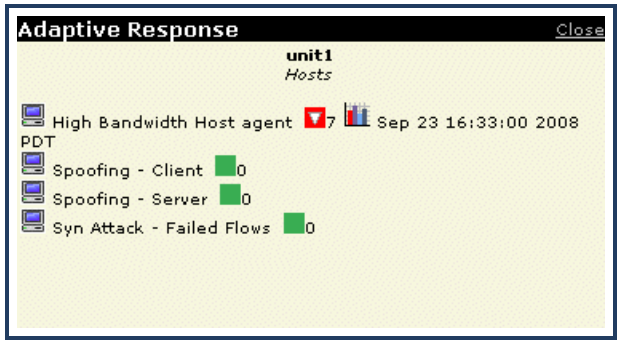

3. In the pop-up window, observe the color of the agent's status indicator.

   Green– During the last evaluation interval, usage of any single host did not exceed the red threshold you defined. That is, no single host used a disproportionate share of the link's bandwidth during the interval.
   Yellow– During the last evaluation interval, usage was approaching the red threshold (yellow is the range between the red and green thresholds).
   Red– During the last evaluation interval, usage of one or more hosts exceeded the defined threshold. In other words, at least one host was using too much bandwidth. If the status indicator is red, you should look at the incident report for details.

   If the agent has ever changed status (from green/yellow to red or from red/yellow to green), a report icon with a link to a detailed report will be available for the agent.

4. If there is a report for the Bandwidth Host agent, click the report icon. The incident report lists all the hosts for the interval, with the host using the most bandwidth at the top of the list. Here is a sample report for the High Bandwidth Host agent.

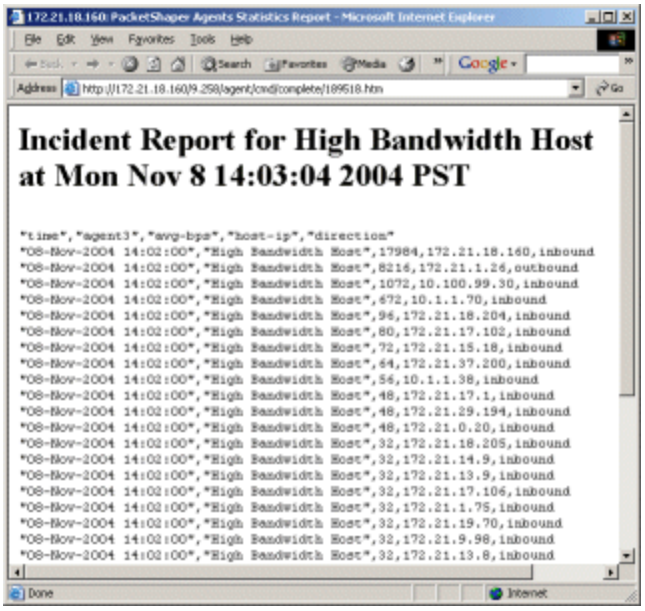

Create an Action File

There are different actions you might want to take if a high bandwidth host is found on your network. Perhaps you just want an email notification that indicates the IP address of the host that is consuming excessive bandwidth. Or perhaps you want to create a class based on the IP address of the bandwidth abuser so that you can track statistics on the host and closely monitor its usage.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and click the **Adaptive Response Setup** link.

2. Click the **edit** button next to the High Bandwidth Host agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, A-Z, -,_, and . (period). Spaces are not allowed. Example: *red-host*

6. In the Contents area, enter the following commands to create an action file for a High Bandwidth Host agent:

```
#Title: red action file for High Bandwidth Host agent
 class new $direction $host-ip outside host:$host-ip
 class rule add $direction/$host-ip inside host:$host-ip
```

```
 send email <address> "High Bandwidth Host found" "Host $host-ip recently used $avg-bps
bps, which is more than the Red Threshold ($RedThreshold percent). Class $direction/$host-
ip was created."
```

where *<address>* is the email address of the recipient. The $host-ip variable represents the IP address of the host that is using the most bandwidth in the interval. The $direction variable represents the host's direction: Inbound or Outbound. The $avg-bps variable is the host's utilization, in bits per second. To summarize, this action file will create a class for the host that is using the most bandwidth; the name of the class will be the host's IP address. You will then be sent an email notifying that these actions have been taken.

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

When the message is sent, the body of the email message would look something like this:

```
Host 65.174.190.201 recently used 762744 bps, which is more than the Red Threshold (10
percent). Class inbound/65.174.190.201 was created.
```

If you simply want to receive an email notification that a host has reached its bandwidth quota, the action file for a Quota Bandwidth Host agent might look like this:

```
#Title: red action file for Quota Bandwidth Host agent
 send email <address> "Host exceeded bandwidth quota" "Host $host-ip used $bytes bytes,
which is more than the defined threshold ($HostUsageThreshold bytes)."
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *red-host.cmd*).

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed.

Use Agents to Detect Spoofing

By using an agent based on the NFPM Side Unknown template, you can detect hosts that may be spoofing on your network. Spoofing attacks send packets that appear to be from a trusted source, but actually contain falsified source and/or destination addresses. The traffic can then gain access to hosts or services that are supposed to be secure.

Two agents based on the NFPM Side Unknown template are created for you automatically: Spoofing - Server and Spoofing - Client. Using these agents, the adaptive response feature will alert you when at least one host has exceeded the new flows per minute threshold. The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when a host is in violation; this method requires that you create an *action file* – a text file of PacketShaper commands. The action file can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to use this feature, the adaptive response feature must be enabled and the Spoofing - Client and Spoofing - Server agents must be turned on.

These agents, as well as most of the other agents in the host category, have two host lists associated with them: exception and violating. In the exception host list, you can add the IP addresses of the hosts you don't want to trigger the agent (such as servers). The violating host list contains the IP addresses of the hosts that exceed the threshold set for the agent; adaptive response automatically adds hosts to this list when they are in violation.

Create an Exception Host List

In order to prevent your servers and other valid hosts from triggering false alarms, you will want to create an exception host list for your Spoofing agents. If you want a single exception host list to apply to all host agents, you can add exception hosts to the default host list that is automatically created for you: exceptionHosts. If you want different host lists for different agents, you will need to create hosts lists with unique names.

1. Access the command-line interface.

2. Type hl show to see the host list names.

3. To create a new host list, use the hl new command. For example:

```
hl new spoofException
```

4. To add hosts to an existing host list, use the hl add command. You can add multiple addresses or subnets by separating each with a space. For example:

```
hl add exceptionHosts 192.21.18.162 192.21.18.165 192.21.18.169
```

5. When modifying the agent (below), make sure to specify this host list name for the ExceptionHosts parameter.

Modify the Agent Settings

By default, the Spoofing agents check every minute to see if a host has exceeded 100,000 new flows per minute. You'll want to set the number of new flows per minute threshold to a value that makes sense for your network. You may also want to change the evaluation interval (for instance, to every 10 minutes).

> **Note:** For the Spoofing agents to work properly, you must set the SideThreshold to a value that is appropriate for your network traffic patterns.

1. On the **Setup** tab, click the **Adaptive Response Setup** link.

2. Make sure the check box in the Agent On column is selected for the Spoofing - Client or Spoofing - Server agent.

3. Click **edit** next to the Spoofing agent you want to modify. The *Edit Agent Entry* window will open.

4. If desired, change the evaluation interval, in minutes (how often the agent should measure its target.)

5. If desired, modify the threshold values:

   SideThreshold– Number of new flows per minute that might indicate a host is spoofing.

   ViolatingHosts – Name of violating host list; adaptive response will automatically create a violating host list using the name supplied here. Any hosts that exceed the agent's SideThreshold will be added to this host list. See Restrict Bandwidth for Violating Hosts for further information.

   ExceptionHosts – Name of exception host list; this should be the name you used in Create an Exception Host List above.

   RedThreshold – The value at which the status of the agent turns red, and a red action file is triggered, if present. For a Spoofing agent, this number represents an unacceptable number of violating hosts; a host is in violation if it exceeds the defined SideThreshold. For instance, if the RedThreshold is set to 1, the agent's status will change to red when at least one host had new flows per minute greater than the defined SideThreshold.

   GreenThreshold– An acceptable number of violating hosts (usually 0).

6. Click **OK** to save your changes.

After creating the agent, you will want to monitor it to see whether any hosts have been in violation.

Check the Agent Status

To check the status of the Spoofing agents:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The Spoofing agents are in the Hosts category.

2. To see the status of each individual agent in the Hosts category, hover your mouse over the Hosts colored status indicator, as shown below.

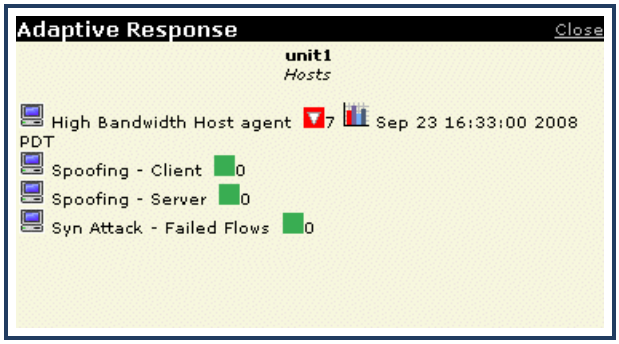

3. In the pop-up window, locate the Spoofing - Client and Spoofing - Server agents. What color are their status indicators?

   Green – During the last evaluation interval, no hosts exceeded the number of new flows per minute you defined for the SideThreshold. (This assumes the GreenThreshold is set to the default value of 0.)
   Red– During the last evaluation interval, at least one host exceeded the SideThreshold. (This assumes the RedThreshold is set to the default value of 1.) If the status indicator is red, you should look at the incident report for details.

   If the agent has ever changed status (from green to red or from red to green, a report icon with a link to a detailed report will be available for that agent.

4. If there is a report for the Spoofing agent, click the report icon. The Spoofing report appears in a separate browser window and lists all the hosts whose new flows per minute were higher than the SideThreshold (during the current interval in addition to previous intervals). These hosts (10.10.10.100 and 10.10.10.101 in the example below) are automatically added to the violatingHosts host list. See Restrict Bandwidth for Violating Hosts for further information about how you can use the violating host list.

Create an Action File that Sends an Email Notification

If you would like to be notified by email when hosts may be spoofing, you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file that sends an email notification:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and choose the **Adaptive Response** setup page.

2. Click the **edit** button next to the Spoofing agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, AZ, -,_, and . (period). Spaces are not allowed. Example: *spoof*

6. In the Contents area, enter the following commands:

```
#Title: red action file for Spoofing - Client agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example:

```
send email raltman@test.com "Spoofing - Client" "Adaptive response has detected at least
```

```
one host with excessive new flows per minute. Look at the incident report for a list of
violating hosts."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

If you like, you can include the defined number of new flows per minute (the $sideThreshold variable) and the number of hosts that were in violation (the $scorevalue variable) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "Spoofing - Client" "Adaptive response has detected $scorevalue
host(s) with more than $sideThreshold new flows per minute. Look at the incident report for
a list of violating hosts."
```

When the message is sent, the body of the email message would look something like this:

```
Adaptive response has detected 2 host(s) with more than 50 new flows per minute. Look at
the incident report for a list of violating hosts.
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *spoof.cmd*).

9. Click **OK**.

An email will be sent if and when the agent's red threshold is crossed during an evaluation interval. If you like, you can create a green action file to notify you when the Spoofing agent returns to a green status.

Restrict Bandwidth for Violating Hosts

Once you know the IP addresses of hosts that might be spoofing, what is your next step? One possibility is to automatically restrict the bandwidth of the violating hosts until you have a chance to research the hosts. To implement this strategy, you create traffic classes that are based on the violating host lists, and then apply appropriate policies and/or partitions to these classes.

1. In Inbound, create a class with a name such as HostViolators.

2. In the **Host List** drop-down list for Inside, select violatingHosts (or the name you specified for the ViolatingHosts parameter).

3. Repeat step 2 for Outside.

4. Click **add class** to create the class.

5. Repeat the above steps to create an Outbound class named HostViolators.

6. Apply appropriate policies and/or partitions to restrict the violating hosts' bandwidth. One strategy might be:

   - Create a partition with size = 0 Kbps, burstable, and a limit of the maximum bandwidth you'd like to devote to all the violating hosts together. For example, you could limit the aggregate of all your violating hosts to 10 percent of your link

size.

- Define a dynamic subpartition giving each host the restricted amount of bandwidth you consider appropriate. For example, you might cap each host at 64 Kbps.

Use Agents to Detect SYN Attacks

A SYN attack is a common denial-of-service (DoS) technique in which an attacker sends multiple SYN packets to a specific host. For each SYN packet received, the target host sends an acknowledgement (SYN-ACK) to the source IP address. Because the target host doesn't receive a response from the attacking machine, it repeatedly attempts to resend the SYN-ACK. When an attacker uses this technique repeatedly, the target host eventually runs out of resources and is unable to handle any more connections, thereby denying service to legitimate users.

When a SYN attack is occurring, there will be hosts with failed flows as well as a high number of new flows per minute. Adaptive response offers several agent templates that monitor flows to/from hosts to help in the detection of SYN attacks:

- **NFPM (New Flows Per Minute) Failed Flow** this agent template detects hosts that have failed flows AND have a high number of new flows per minute.
- **Failed Flow Ratio** this agent template detects hosts that have a high ratio of failed flows compared to new client flows per minute.
- **Host Info Variables** this agent template can monitor several different types of flows, one of which is failed flows. An agent, called Syn Attack - Failed Flow, is automatically created for you; you can modify the default settings if necessary.

After you have set up one or more of these agents, the adaptive response feature will alert you when a possible SYN attack is in process. The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent to you; this method requires that you create an *action file*– a text file of PacketShaper commands. The action file can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to use this feature, the adaptive response feature must be enabled. If you want to use the preconfigured SYN Attack agent but want to review or modify the settings, refer to Modify the Syn Attack - Failed Flow Agent Settings, below. If you want to create your own agent based on either the NFPM Failed Flow or Failed Flow Ratio template, see Add a New Adaptive Response Agent for instructions on creating a new agent.

> **Note:** These agents have two host lists associated with them: exception and violating. In the exception host list, you can add the IP addresses of the hosts you don't want to trigger the agent (such as servers). The violating host list contains the IP addresses of the hosts that exceed the threshold set for the agent; adaptive response automatically adds hosts to this list when they are in violation.

Create an Exception Host List

In order to prevent your servers and other valid hosts from triggering false alarms, you will want to create an exception host list for your agents. If you want a single exception host list to apply to all host agents, you can add exception hosts to the default host list that is automatically created for you: exceptionHosts. If you want different host lists for different agents, you will need to create hosts lists with unique names.

1. Access the command-line interface.

2. Type hl show to see the host list names.

3. To create a new host list, use the hl new command. For example:

```
hl new synException
```

4. To add hosts to an existing host list, use the hl add command. You can add multiple addresses or subnets by separating each with a space. For example:

```
hl add exceptionHosts 192.21.18.162 192.21.18.165 192.21.18.169
```

5. When modifying the agent (below), make sure to specify this host list name for the ExceptionHosts parameter.

Modify the Syn Attack - Failed Flow Agent Settings

By default, the Syn Attack - Failed Flow agent checks every minute to see if a host has exceeded 100,000 failed flows per minute. You'll want to set the number of failed flows threshold to a value that makes sense for your network. You may also want to change the evaluation interval (for instance, to every 10 minutes).

> **Note:** For the agent to work properly, you must set the FlowsThreshold to a value that is appropriate for your network traffic patterns.

1. On the **Setup** tab, choose the **Adaptive Response** setup page.
2. Make sure the check box in the Agent On column is selected for the Syn Attack - Failed Flow agent.
3. Click **edit** next to the agent. The *Edit Agent Entry* window will open.
4. If desired, change the evaluation interval, in minutes (how often the agent should measure its target.)
5. If desired, modify the threshold values:

   FlowsThreshold– Number of failed flows per minute that might indicate a host is being attacked.

   Side – Traffic direction to be monitored (both, inside, outside)

   ViolatingHosts– Name of violating host list; adaptive response will automatically create a violating host list using the name supplied here. Any hosts that exceed the agent's FlowsThreshold will be added to this host list. See Restrict Bandwidth for Violating Hosts for further information.

   ExceptionHosts– Name of exception host list; this should be the name you used in Create an Exception Host List above.

   RedThreshold –The value at which the status of the agent turns red, and a red action file is triggered, if present. For a Syn Attack agent, this number represents an unacceptable number of violating hosts; a host is in violation if it exceeds

the defined FlowsThreshold. For instance, if the RedThreshold is set to 1, the agent's status will change to red when at least one host had failed flows per minute greater than the defined FlowsThreshold.

GreenThreshold– An acceptable number of violating hosts (usually 0).

6. Click **OK** to save your changes.

After creating the agent, you will want to monitor it to see whether any hosts have been in violation.

Check the Agent Status

To check the status of the SYN attack agent(s):

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The SYN attack agents are in the Hosts category.

2. To see the status of each individual agent in the Hosts category, hover your mouse over the Hosts colored status indicator, as shown below.

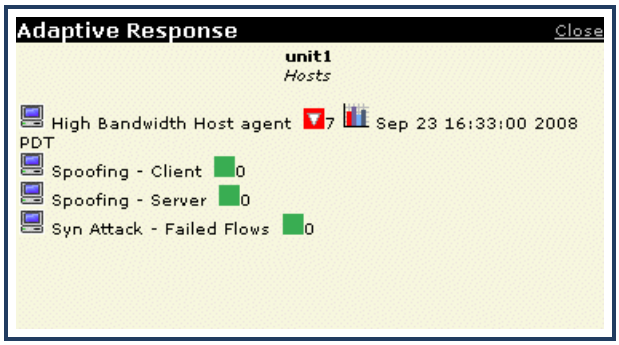

3. In the pop-up window, locate your SYN attack agent(s). What color are their status indicators?

   Green– During the last evaluation interval, no hosts exceeded the number of flows per minute you defined for the FlowsThreshold (for agents based on the Host Info Variables or NFPM - Failed Flow templates) or RatioThreshold (for agents based on the Failed Flow Ratio template). Note: This assumes the GreenThreshold is set to the default value of 0.
   Red – During the last evaluation interval, at least one host exceeded the FlowsThreshold (or RatioThreshold). Note: This assumes the RedThreshold is set to the default value of 1.

   If the status indicator is red, you should look at the incident report for details.

4. If there is a report for the SYN attack agent, click the report icon. The report appears in a separate browser window and lists all the hosts who exceeded the defined thresholds (during the current interval in addition to previous intervals). These hosts (10.10.10.100 in the example below) are automatically added to the violatingHosts host list. See Restrict Bandwidth for Violating Hosts for further information about how you can use the violating host list.

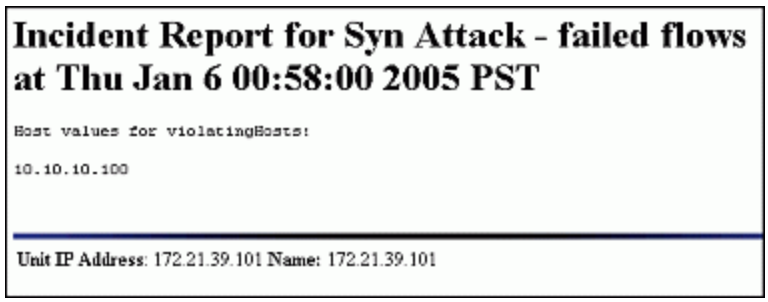

Create an Action File that Sends an Email Notification

If you would like to be notified by email when a SYN attack might be occurring, you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file that sends an email notification:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and choose the **Adaptive Response** setup page.

2. Click the **edit** button next to the SYN attack agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, AZ, -,_, and . (period). Spaces are not allowed. Example: *synattak*

6. In the Contents area, enter the following commands:

```
#Title: red action file for SYN Attack agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example:

```
send email raltman@test.com "Syn Attack may be in process" "Adaptive response has detected
at least one host that may be involved in a Syn Attack. Look at the incident report for a
list of violating hosts."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email

address.

If you like, you can include the defined number of flows per minute (the $flowsThreshold or $ratioThreshold variable) and the number of hosts that were in violation (the $scorevalue variable) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "Syn Attack may be in process" "Adaptive response has detected $scorevalue host(s) with more than $flowsThreshold failed flows per minute. Look at the incident report for a list of violating hosts."
```

When the message is sent, the body of the email message would look something like this:

```
Adaptive response has detected 2 host(s) with more than 100 failed flows per minute. Look at the incident report for a list of violating hosts.
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *synattak.cmd*).

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed. If you like, you can create a green action file to notify you when the agent returns to a green status.

Restrict Bandwidth for Violating Hosts

Once you know the IP addresses of hosts that might be involved in a SYN attack, what is your next step? One possibility is to automatically restrict the bandwidth of the violating hosts until you have a chance to research the hosts. To implement this strategy, you create traffic classes that are based on the violating host list, and then apply appropriate policies and/or partitions to these classes.

1. In Inbound, create a class with a name such as HostViolators.

2. In the **Host List** drop-down list for Inside, select violatingHosts (or the name you specified for the ViolatingHosts parameter).

3. Repeat step 2 for Outside.

4. Click **add class** to create the class.

5. Repeat the above steps to create an Outbound class named HostViolators.

6. Apply appropriate policies and/or partitions to restrict the violating hosts' bandwidth. One strategy might be:

- Create a partition with size = 0 Kbps, burstable, and a limit of the maximum bandwidth you'd like to devote to all the violating hosts together. For example, you could limit the aggregate of all your violating hosts to 10 percent of your link size.

- Define a dynamic subpartition giving each host the restricted amount of bandwidth you consider appropriate. For example, you might cap each host at 64 Kbps.

Create an Agent that Monitors a Partition Variable

The Partition Utilization agent monitors a partition's average rate over a time interval. But what if you want to monitor another aspect of a partition, such as the number of dropped packets or the amount of time (or percent of time) during which allocated bandwidth exceeded the partition minimum size? This is a job for a User Event Emulation agent based on a partition ME variable.

In this example, we will create a new agent called Partition Over Limit. This agent will monitor the partition-over-limit-time% measurement variable (the percentage of time that allocated bandwidth in a class exceeded the partition minimum size). When the class' traffic exceeds the partition minimum size more than a certain percentage of time, the adaptive response feature will alert you that a threshold has been exceeded. The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when the red threshold is crossed; this method requires that you create an *action file* – a text file of PacketShaper commands. The action file can take other actions, as well, such as modifying the partition size, and can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to create agents, the adaptive response feature must be enabled.

Create a New Agent

To create an agent for a partition variable:

1. On the **Setup** tab, choose the **Adaptive Response** setup page.
2. Click **add**. The *Add Agent Entry* window will open.
3. Click the **create a new agent from this agent template** drop-down list, and choose *Partition ME Variables*.
4. Enter a unique identifying name for the new agent in the **Name** field, for example *FTP Partition Over Limit*.
5. Click **OK and Edit Parms**. This will open the *Edit Agent Entry* window where you can define the parameter values for your new agent.
6. In the **ClassName** field, enter the name of the class to be monitored (for example, *Inbound/FTP*). Note that the class must already have a partition defined.
7. In the **MeVariableName** field, enter the measurement variable to be monitored (*partition-over-limit-time%*). For a complete list of measurement variables you can use, see Measurement Variables; you can select any variable in the Common and Partition categories.

8. Specify the threshold values:

   RedThreshold – The value at which the status of the agent turns red, and a red action file is triggered, if present. For a Partition Over Limit agent, this number represents the percent of time at which bandwidth allocation exceeded the minimum partition size. For instance, if the RedThreshold is set to 60, the agent's status will change to red when the class' bandwidth allocation exceeded the minimum partition size more than 60 percent of the time in the MeDuration (see step 11 below).
   GreenThreshold – The value at which the status of the agent turns green, and a green action file is triggered, if present. For instance, if the GreenThreshold is set to 30, the agent's status will be green when the bandwidth allocation of the class exceeds the minimum partition size less than 30 percent of the time in the MeDuration.

9. If desired, change the values for **GreenOperator** and/or **RedOperator**.

   In this example, the values for GreenOperator (LT less than) and RedOperator (GT greater than) are appropriate. This agent will have a red status when the current value of the measurement variable is greater than (GT) the RedThreshold value and a green status when the current value is less than (LT) the GreenThreshold value.

10. If desired, change the **evaluation interval**, in minutes (how often the agent should measure its target).

11. Specify the **MeDuration**, the period of time over which measurements will be looked at.

    The MeDuration can be different from the evaluation interval. For example, suppose you set the evaluation interval to 5 minutes and the MeDuration to 60 minutes. The agent will look at the last hour of measurement data every five minutes to see if the red threshold has been exceeded. At 12:00, the agent would analyze data collected between 11:00 and 12:00; at 12:05, the agent would analyze data collected between 11:05 and 12:05; and so forth.

12. Click **OK** to save your changes.

After creating the agent, you will want to monitor it to see how often traffic in the class exceeded its partition minimum size.

Check the Agent Status

To check the status of the Partition Over Limit agent:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The Partition Over Limit agent is in the User Event Emulation category (the last one).

2. To see the status of each individual agent in the User Event Emulation category, hover your mouse over the colored status indicator, as shown below.

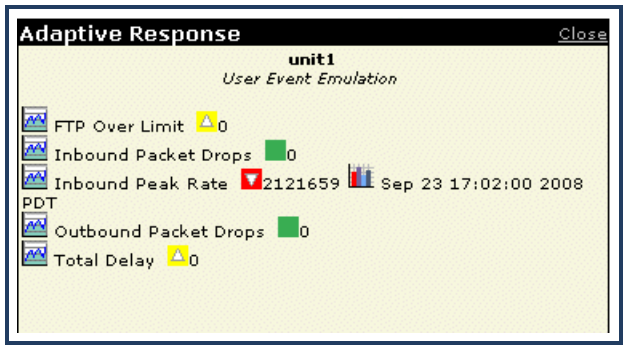

3. In the pop-up window, locate the Partition Over Limit agent that you defined. The value (16 in the figure above) indicates the percent of time that bandwidth allocation in the class exceeded the partition minimum size during the MeDuration.

   What color is its status indicator?

   Green – During the MeDuration, traffic in the partition did not exceed its minimum size very often (less than the green threshold).
   Yellow– During the MeDuration, the score (the value of the measurement variable) was between the red and green thresholds.
   Red – During the MeDuration, traffic in the class exceeded its partition minimum size frequently (more than the red threshold). If the status indicator is red, you should look at the incident report for details.

   If the agent has ever changed status (from green/yellow to red or from red/yellow to green, a report icon with a link to a detailed report will be available for that agent.

4. If there is a report for the Partition Over Limit agent, click the report icon. The incident report appears in a separate browser window.

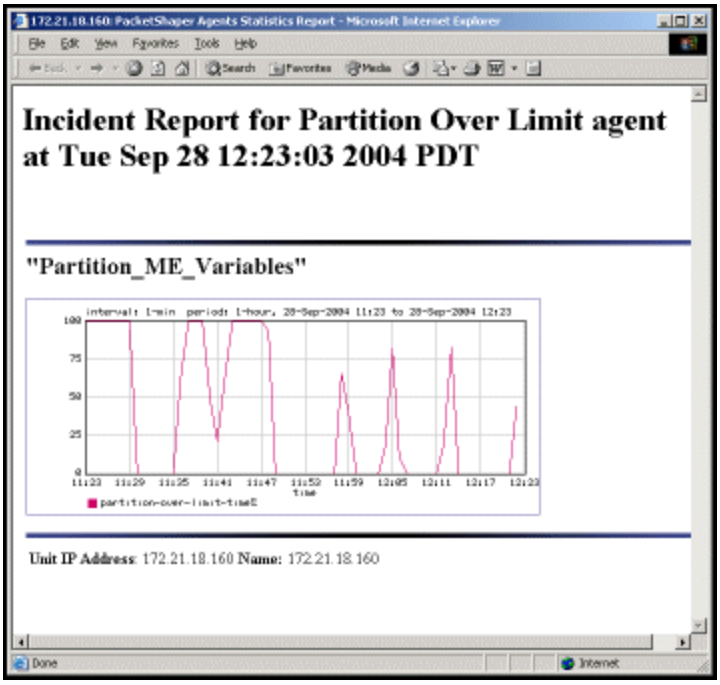

Create an Action File that Sends an Email Notification

If you would like to be notified by email when the class exceeds its partition minimum size too frequently (with "too frequently" being defined as the red threshold), you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file that sends an email notification:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and click the **Adaptive Response Setup** link.

2. Click the **edit** button next to the Partition Over Limit agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, AZ, -,_, and . (period). Spaces are not allowed. Example: *partover*

6. In the Contents area, enter the following commands:

```
#Title: red action file for Partition Over Limit agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example:

```
send email raltman@test.com "FTP partition" "Bandwidth allocation in the Inbound/FTP class
frequently exceeded its partition minimum."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

Alternatively, you can include the variables for the agent name, ($agentname), class name ($class-id), value of the measurement variable ($me-value), and the red threshold ($RedThreshold) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "$agentname agent" "The $class-id class was over its partition
minimum size $me-value percent of the time. This is over the red threshold
($RedThreshold)."
```

When the message is sent, the body of the email message would look something like this:

```
The Inbound/FTP class was over its partition minimum size 67 percent of the time. This is
over the red threshold (60).
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *partover.cmd*).

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed. If you like, you can create a green action file to notify you when the Partition Over Limit agent returns to a green status.

Create an Agent that Monitors a Class Variable

You can create an adaptive response agent to monitor any class measurement variable available in PacketShaper. For example, you can monitor application availability (app-availability%), average round trip time (avg-round-trip-time), network delay (network-delay-avg), or packet exchange time (pkt-exchange-time-avg) of a traffic class. See Measurement Variables for a complete list.

In this example, we will create a new agent called Total Delay. This agent will monitor the total-delay-avg measurement variable (average number of milliseconds to complete the class' transactions; includes network delay and server delay) for a specific traffic class. When the class' total delay exceeds a predefined threshold, the adaptive response feature will alert you that a threshold has been exceeded. The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when the red threshold is crossed; this method requires that you create an *action file*– a text file of PacketShaper commands. The action file can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to create agents, the adaptive response feature must be enabled.

Create a New Agent

To create an agent for a class measurement variable:

1. On the **Setup** tab, select the **Adaptive Response Setup** setup page.
2. Click **add**. The *Add Agent Entry* window opens.
3. Click the **create a new agent from this agent template** drop-down list, and choose *Class ME Variables*.
4. Enter a unique identifying name for the new agent in the **Name** field, for example *Citrix Delay*.
5. Click **OK and Edit Parms**. This will open the *Edit Agent Entry* window where you can define the parameter values for your new agent.
6. In the **ClassName** field, enter the name of the class to be monitored (for example, *Inbound/Citrix*).
7. In the **MeVariableName** field, enter the measurement variable to be monitored (*total-delay-avg*). For a complete list of measurement variables you can use, see Measurement Variables; you can select any variable in the Common and Class categories.
8. Specify the threshold values:

RedThreshold– The value at which the status of the agent turns red, and a red action file is triggered, if present. For a Total Delay agent, this number represents the average delay in milliseconds. For instance, if the RedThreshold is set to 500, the agent's status will change to red when the class' average network and server delay is over 500 ms in the MeDuration (see step 11 below).

GreenThreshold– The value at which the status of the agent turns green, and a green action file is triggered, if present. For instance, if the GreenThreshold is set to 300, the agent's status will be green when the average delay is less than 300 ms in the MeDuration.

9. If desired, change the values for **GreenOperator** and/or **RedOperator**.

   In this example, the values for GreenOperator (LT– less than) and RedOperator (GT – greater than) are appropriate. This agent will have a red status when the current value of the measurement variable is greater than (GT) the RedThreshold value and a green status when the current value is less than (LT) the GreenThreshold value. You may want to change the values to GTE (greater than or equal to) and LTE (less than or equal to).

10. If desired, change the **evaluation interval**, in minutes (how often the agent should measure its target.)

11. Specify the **MeDuration**, the period of time over which measurements will be looked at.

    The MeDuration can be different from the evaluation interval. For example, suppose you set the evaluation interval to 5 minutes and the MeDuration to 60 minutes. The agent will look at the last hour of measurement data every five minutes to see if the red threshold has been exceeded. At 12:00, the agent would analyze data collected between 11:00 and 12:00; at 12:05, the agent would analyze data collected between 11:05 and 12:05; and so forth.

12. Click **OK** to save your changes.

After creating the agent, you will want to monitor it to see if the average network and server delay in the class was excessive.

Check the Agent Status

To check the status of the Total Delay agent:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The Total Delay agent is in the User Event Emulation category (the last one).

2. To see the status of each individual agent in the User Event Emulation category, hover your mouse over the colored status indicator, as shown below.

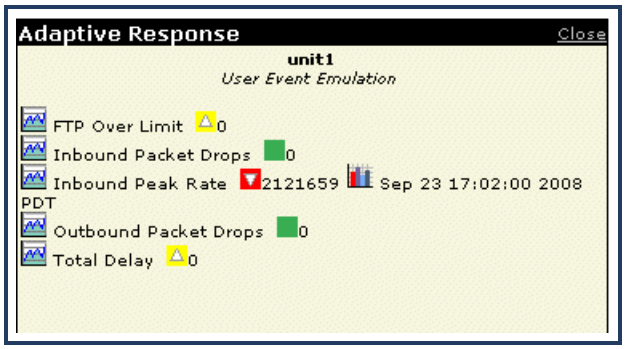

3. In the pop-up window, locate the Total Delay agent that you defined. The value (for example, 1970) indicates the number of milliseconds of the class' average network and server delay during the MeDuration.

   What color is its status indicator?

   Green– During the MeDuration, delay in the class was at an acceptable level (less than the green threshold).
   Yellow– During the MeDuration, the score (the value of the measurement variable) was between the red and green thresholds.
   Red– During the MeDuration, delay in the class was at an unacceptable level (more than the red threshold). If the status indicator is red, you should look at the incident report for details.

   If the agent has ever changed status (from green/yellow to red or from red/yellow to green, a report icon with a link to a detailed report will be available for that agent.

4. If there is a report for the Total Delay agent, click the report icon. The incident report appears in a separate browser window.

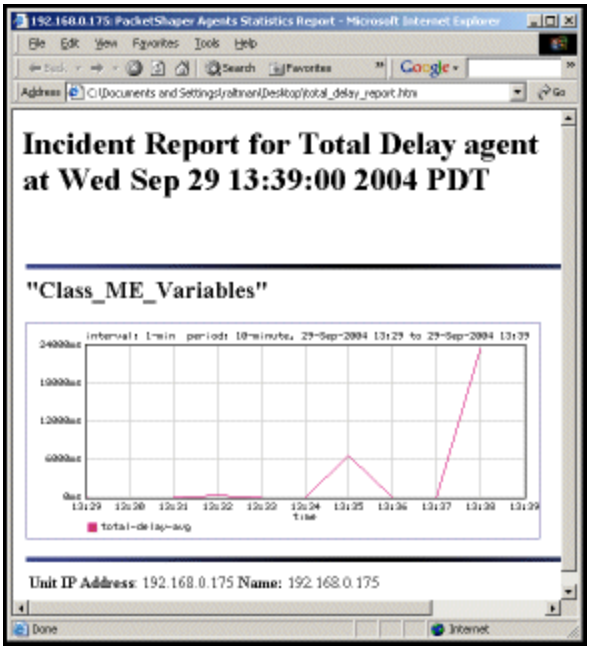

Create an Action File that Sends an Email Notification

If you would like to be notified by email when a class' delay is excessive (that is, above the red threshold), you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file that sends an email notification:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and select the **Adaptive Response** setup page.

2. Click the **edit** button next to the Total Delay agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, AZ, -,_, and . (period). Spaces are not allowed. Example: *appdelay*

6. In the Contents area, enter the following commands:

```
#Title: red action file for Total Delay agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example:

```
send email raltman@test.com "Citrix is too slow" "The average delay (server plus network)
for Inbound/Citrix was at an unacceptable level."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

Alternatively, you can include the variables for the agent name ($agentname), class name ($class-id), value of the measurement variable ($me-value), and the red threshold ($RedThreshold) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "$agentname agent" "The average delay (server plus network) of
the $class-id class was $me-value ms. This is over the red threshold ($RedThreshold ms).
Go to the Info tab to view the report for the $agentname agent."
```

When the message is sent, the body of the email message would look something like this:

```
The average delay (server plus network) of the Inbound/Citrix class was 70 ms. This is over
the red threshold (60 ms). Go to the Info tab to view the report for the Citrix Delay
agent.
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *appdelay.cmd*).

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed. If you like, you can create a green action file to notify you when the Total Delay agent returns to a green status.

Create an Agent that Monitors a Link Variable

The Network Health category has several agents that monitor your Inbound and Outbound links (Traffic Performance and Partition Utilization), but if you want to monitor other types of statistics on your link, you can create an agent that analyzes any link measurement variable available in PacketShaper. For example, you can monitor a link's peak rate (peak-bps), percentage of retransmitted TCP packets (tcp-retx-pkts%), or percentage of refused connections (tcp-conn-server-refuses%). See Measurement Variables for a complete list.

In this example, we will create a new agent called Inbound Peak Rate. This agent will monitor the peak-bps measurement variable for the Inbound link. When the class' peak rate exceeds a user-defined threshold, the adaptive response feature will alert you that a threshold has been exceeded. The standard alert is a red indicator on the Info tab (PacketShaper) or Configuration tab (PolicyCenter). In addition, you can have an email, SNMP trap, or syslog message sent when the red threshold is crossed; this method requires that you create an *action file*– a text file of PacketShaper commands. The action file can include any CLI command that can run without additional user input. See "Adaptive Response Action Files Overview" on page 58 for more information.

In order to create agents, the adaptive response feature must be enabled.

Create a New Agent

To create an agent for a link measurement variable:

1. On the **Setup** tab, click the **Adaptive Response Setup** link.
2. Click **add**. The *Add Agent Entry* window will open.
3. Click the **create a new agent from this agent template** drop-down list, and choose *Link ME Variables*.
4. Enter a unique identifying name for the new agent in the **Name** field, for example *Inbound Peak Rate*.
5. Click **OK and Edit Parms**. This will open the *Edit Agent Entry* window where you can define the parameter values for your new agent.
6. In the **ClassName** field, enter the name of the link to be monitored (for example, */Inbound*).
7. In the **MeVariableName** field, enter the measurement variable to be monitored (*peak-bps*). For a complete list of measurement variables you can use, see Measurement Variables; you can select any variable in the Common and Link categories.

8. Specify the threshold values:

   RedThreshold– The value at which the status of the agent turns red, and a red action file is triggered, if present. For a Peak Rate agent, this number represents the highest rate, in bits per second, recorded during the MeDuration (see step 11 below). For instance, if the RedThreshold is set to 500000, the agent's status will change to red when the link's peak rate is over 500 kbps in the MeDuration.
   GreenThreshold– The value at which the status of the agent turns green, and a green action file is triggered, if present. For instance, if the GreenThreshold is set to 300000, the agent's status will be green when the peak rate is less than 300 kbps in the MeDuration.

9. If desired, change the values for **GreenOperator** and/or **RedOperator**.

   In this example, the values for GreenOperator (LT – less than) and RedOperator (GT – greater than) are appropriate. This agent will have a red status when the current value of the measurement variable is greater than (GT) the RedThreshold value and a green status when the current value is less than (LT) the GreenThreshold value.

10. If desired, change the **evaluation interval**, in minutes (how often the agent should measure its target).

11. Specify the **MeDuration**, the period of time over which measurements will be looked at.

    The MeDuration can be different from the evaluation interval. For example, suppose you set the evaluation interval to 5 minutes and the MeDuration to 60 minutes. The agent will look at the last hour of measurement data every five minutes to see if the red threshold has been exceeded. At 12:00, the agent would analyze data collected between 11:00 and 12:00; at 12:05, the agent would analyze data collected between 11:05 and 12:05; and so forth.

12. Click **OK** to save your changes.

After creating the agent, you will want to monitor it to see if the link's peak rate surpassed the red threshold.

Check the Agent Status

To check the status of the Peak Rate agent:

1. Click the **Info** tab. The PacketShaper s Info tab shows an icon for each agent category, and a colored status indicator for each. The Peak Rate agent is in the User Event Emulation category (the last one).

2. To see the status of each individual agent in the User Event Emulation category, hover your mouse over the colored status indicator, as shown below.

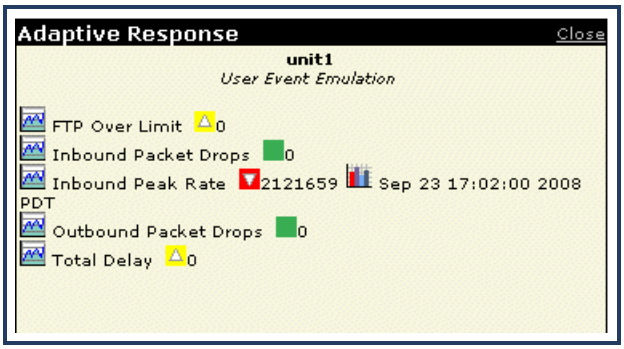

3. In the pop-up window, locate the Peak Rate agent that you defined. The value (for example, 903848) indicates the link's peak rate recorded during the MeDuration.

   What color is its status indicator?

   Green– During the MeDuration, the peak on the link was at an acceptable rate (less than the green threshold).
   Yellow– During the MeDuration, the score (the value of the measurement variable) was between the red and green thresholds.
   Red– During the MeDuration, the peak on the link was at an unacceptable rate (more than the red threshold). If the status indicator is red, you should look at the incident report for details.

   If the agent has ever changed status (from green/yellow to red or from red/yellow to green, a report icon with a link to a detailed report will be available for that agent.

4. If there is a report for the Peak Rate agent, click the report icon. The incident report appears in a separate browser window.

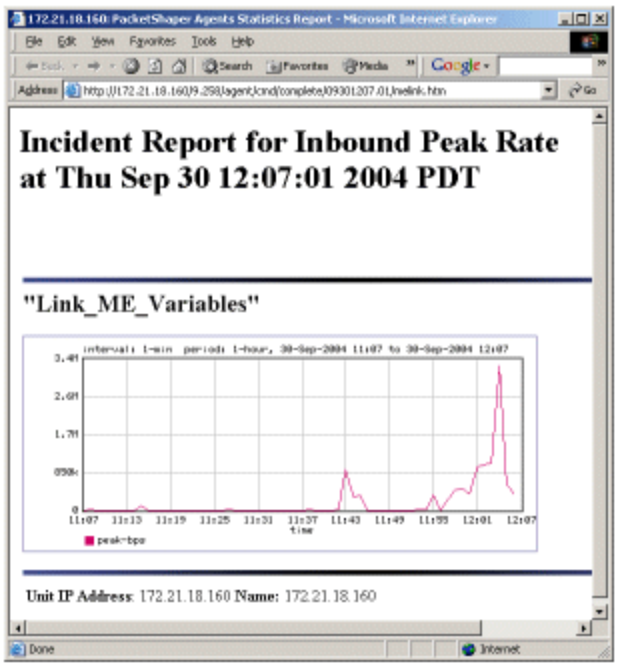

Create an Action File that Sends an Email Notification

If you would like to be notified by email when the peak rate on a link is too high (that is, above the red threshold), you can create a command file that contains a single line: the **send email** command. By designating this command file as the red action file, an email will be sent when the red threshold is crossed.

> **Note:** To send email notification, you must configure an SMTP server in PacketShaper.

To create an action file that sends an email notification:

1. If the Info tab is currently displayed, click the **Settings** link; otherwise, go to the **Setup** tab and choose the **Adaptive Response** setup page.

2. Click the **edit** button next to the Peak Rate agent.

3. Click the **browse files** button. The File Browser window opens and shows the contents of the 9.258/agent/cmd directory (where action files need to be stored).

4. Click the **new cmd file** button. A command file window opens.

5. In the **File Name** field, enter a unique name for your action file up to eight characters long, including 0-9, a-z, AZ, -,_, and . (period). Spaces are not allowed. Example: *peak*

6. In the Contents area, enter the following commands:

```
#Title: red action file for Peak Rate agent
 send email <address> "<subject>" ["<body>"]
```

where *<address>* is the email address of the recipient. For example:

```
send email raltman@test.com "Inbound peak rate" "The peak rate on the Inbound link was
above the red threshold."
```

Note: If you want to send an email notification to more than one recipient, repeat the **send email** command for each email address.

Alternatively, you can include the variables for the agent name ($agentname), class name ($class-id), value of the measurement variable ($me-value), and the red threshold ($RedThreshold) in the *<subject>* or *<body>*, as the following example shows:

```
send email raltman@test.com "$agentname agent" "The peak rate of the $class-id link was
$me-value bps. This is over the red threshold ($RedThreshold BPS). Go to the Info tab to
view the report for the $agentname agent."
```

When the message is sent, the body of the email message would look something like this:

```
The peak rate of the /Inbound link was 982438 BPS This is over the red threshold (900000
```

```
BPS). Go to the Info tab to view the report for the Inbound Peak Rate agent."
```

7. Click **save**.

8. Enter the name of your action file in the **Red Action File** field (for example, *peak.cmd)*.

9. Click **OK**.

After the next evaluation interval, an email will be sent if and when the agent's red threshold is crossed. If you like, you can create a green action file to notify you when the Peak Rate agent returns to a green status.

Event Notification Tasks

A PacketShaper event is a network situation that can trigger an alert message sent via email, an SNMP trap message, and/or a Syslog message. Network and system administrators rely on event notification to proactively manage network resources. For example, you can use event notification to automatically report to you when there are unacceptable network delays or when there isn't enough bandwidth to satisfy guaranteed rate demand. By having PacketShaper send you notifications automatically, you have the opportunity to be aware of situations before they trigger user complaints. For more information, see the Overview of Event Notification.

This section includes tasks for configuring the event notification feature, defining and registering events, and viewing the event status report.

Set Up Event Notification

The PacketShaper software can notify you when an event occurs via email, Simple Network Management Protocol (SNMP) traps, or Syslog. While registering an event, you select which method(s) you want to use for notification of that event. However, before you can do this, you need to set up your system to handle these functions. For more information about the PacketShaper event feature, see Overview of Event Notification and Notify Someone of Situations of Interest.

Defining the SNMP Trap Destination

In order to send SNMP traps when an event occurs, PacketShaper needs to know where to send the traps—the SNMP trap destination. You can define up to eight IP addresses in the Setup window's SNMP page. For details, see Configure PacketShaper for SNMP Support.

Setting Up Email Notification

PacketShaper uses the SMTP protocol for email messaging. To configure email notification, you need to:

1. Specify the SMTP server location and the sender's email account.

2. Indicate who will receive the email notification of events when thresholds are crossed.

3. Verify the email configuration by sending a test message.

See Set Up Email Notification for details.

Setting Up Syslog on a PacketShaper

The logging feature gives administrators a way to centrally log and analyze configuration events and system error messages. To set up logging, you need to define the Syslog servers and then enable the logging feature. See Set Up Syslog for details.

Set Up Email Notification

PacketShaper uses the Simple Mail Transfer Protocol (SMTP) for email messaging. To configure email notification, you need to:

1. Configure an SMTP server.

2. Define email recipients (those who will receive the email notification of events when thresholds are crossed).

Define Email Recipients

Add recipients for event email by using the event email command from the CLI, as described below.

To configure email recipients:

1. At the command-line prompt, type **event email add** and press Enter.

2. When prompted for the recipient name, type the email name and/or address.

To send to more than one recipient, simply leave a space between each email address. You can list up to four recipient addresses.

3. To see a list of event email recipients, type **event show** and press Enter. The recipient list appears at the beginning of the output (above the event list).

To remove a recipient, use the **event email delete *<recipient>*** command, where *<recipient>* is the name or email address of the recipient you want to delete.

To verify that the event email configuration is valid, use the email test *<recipient>* command from the CLI. To verify the user event recipient list is working, use the event test-email command.

Use Predefined Events

PacketShaper includes a set of predefined events to facilitate event configuration. If a predefined event already exists, you can skip the definition process and go directly to registering the event.

PacketShaper Predefined Event List

The table below lists each predefined event name, the default frequency, and the event expression. The event expression includes the name of the measurement variable (such as network-delay-avg), the object type (class, partition, or link), and the relational operator (> or <). The *$1* represents the threshold variable (when you register the event, the $1 will be replaced with the threshold value you specify).

Predefined Event Name	Frequency (Default)	Event Expression
AppDelayNetworkInduced	1 min	network-delay-avg.class>$1
AppDelayServerInduced	1 min	server-delay-avg.class>$1
AppDelayTotal	1 min	network-delay-avg.class>$1
ClassAverageBitRateHigh	1 min	avg-bps.class>$1
ClassAveTransSize	1 min	trans-bytes-avg.class>$1
ClassConnAbortsHigh	1 min	tcp-conn-aborts%.class>$1
ClassPeakRateHigh	1 hr	peak-bps.class>$1
ClassStarvation	1 min	guar-rate-fails.class>$1
LicenseOverflows	1 min	license-overflows.class>$1
LinkAverageBitRateHigh	1 min	avg-bps.link>$1
LinkConnAbortsHigh	1 min	tcp-conn-aborts%.link>$1
LinkPeakRateHigh	1 hr	peak-bps.link>$1
NetworkPathInefficiency	1 min	tcp-efficiency%.link<$1
NormalizedNetworkDelay	1 min	normalized-network-delay-avg.class>$1
PartitionAverageBitRateHigh	1 min	avg-bps.partition>$1
PartitionConnAbortsHigh	1 min	tcp-conn-aborts%.partition>$1
PartitionOverLimit	1 min	partition-over-limit-time%>$1
PartitionPeakRateHigh	1 hr	peak-bps.partition>$1
PartitionStarvation	1 min	lowest-fully-satisfied-priority.partition>$1
ServiceLevelError	1 min	service-level-errors.class>$1
SlowTransactions	1 min	service-level%.class>$1
UnsolicitedICMP	1 min	unsolicited-icmp.link>$1

Define an Event

When you define an event, you build an *expression* – the condition for which you want to be notified. An expression includes the measurement variable, type of object the event is measuring (link, class, or partition), and a relational operator (such as >). Here's an example of an expression you could build:

```
network-delay-avg.class>$1
```

This expression looks at the network-delay-avg measurement variable for a class and checks whether it is greater than a threshold variable (represented by *$1*). The class name and the threshold value are not set when you define the event – they are set when you register it.

To define a new event:

1. Access the CLI.

2. Type **event new** and press Enter.

3. Enter a descriptive name for the event (spaces are not allowed).

4. Specify the type of object the event is measuring: enter **l** for link, **p** for partition, or **c** for class.

5. Type the name of the measurement variable. Be sure to type the variable exactly as it's listed in the table.

6. Enter the default checking interval (for example, *1m* for one minute, *1h* for one hour).

Note: When setting the notification-checking interval, take into account the sampling interval of the measurement variable. Variables are sampled on one-minute and one-hour intervals so valid checking intervals are **1m** or **1h**.

7. Enter a relational operator.

You can define a maximum of 32 events, in addition to the predefined events included with PacketShaper.

Once you have defined an event, you can register it. When you register an event, you supply a threshold value, which serves as the trigger for event notification. See Register an Event.

Register an Event

After you have defined an event or selected a predefined event, you can initiate event-checking and notification by registering the event. When you define an event, you specify a measurement variable in a condition you want to check. Then, when you register the event, you cite the particular link, partition, or class that should be checked for the threshold condition, as well as the specific threshold value. Up to 32 events can be registered at one time.

When you register an event, you specify the following values:

- The name of the event's object (the path to the link, partition, or class)
- Alarm threshold – the value that triggers event notification
- Re-arm level – the value that tells PacketShaper that it's okay to once again send notification.
- Notification type (email, SNMP trap, and/or Syslog)
- Interval for event checking, if other than the default for the measurement variable
- Maximum number of event notifications in a day

To register a user-defined or predefined event:

1. Access the CLI.

2. Type **event register** and press Enter. A list of available events appears.

3. Referring to the list on the screen, enter the name of the event to be registered. (You don't need to type the entire name just type enough to make it unique.)

4. Enter the path to the object you are testing. For example, if you are testing a link, type **outbound** or **inbound**. If you are testing a class or partition, type the path to the traffic class (such as */Inbound/HTTP*).

5. Enter a value for the alarm threshold.

6. Enter a re-arm value.

Notes:

- The event-registration process requires that you set a re-arm level to prevent over-notification of a condition. Fluctuating network activity could cause multiple occurrences of the same event, and you usually do not want to be alerted each and every time. For example, you could configure an event to notify you when your web server's tcp-conn-aborts% exceeds 70%. This condition could occur many times within a short time frame, but all of these occurrences signal only one event. In this example, you might set the re-arm value to 60% to tell PacketShaper to "notify me after the aborts drop back down to a normal level, and crosses the threshold *again*." The re-arm level clears the condition so that you'll only be notified of truly new events.
- PacketShaper sends a notification when the re-arm value is crossed, allowing you to be alerted automatically when the event has been cleared.

7. Press Enter to specify the default checking interval, or type a different value.

Valid values are **1m** or **1h** (for PacketShaper ISP, **1m** or **4h**).

8. Indicate whether you want an email sent when an event occurs: type **y** for yes or **n** for no.

> **Note:** This feature requires that the system is configured for email notification. See Set Up Email Notification.

9. Indicate whether you want an SNMP trap sent when an event occurs: type **y** for yes or **n** for no.

> **Note:** This feature requires that you have defined the SNMP trap destination. See " Configure SNMP on PacketShaper" on page 334.

10. Indicate whether you want to send an event message to the Syslog server: type **y** for yes or **n** for no.

> **Note:** This feature requires that you have defined a Syslog server. See Set Up Syslog.

11. Specify the maximum number of event notifications to be sent in a 24-hour period.

> **Note:** If the daily limit of event notifications is reached, you will not be notified of the events, and the events will not be recorded into the log file.

After the last prompt, you are provided with an event registration identifier. For example:

```
Event registration id: 1
```

You can use this number later if you want to unregister the event.

When an event exceeds the predefined threshold value, the event is in violation and PacketShaper will automatically send out notification. PacketShaper will also send a notification when the re-arm value is crossed, allowing you to be alerted automatically when the event has been cleared.

To display a list of your registered events and determine which are compliant, disarmed, or in violation, see View Event Status.

View Event Status

On the *Monitor Events* screen in the browser interface, you can check the status of the events you have registered, such as whether the threshold value has been crossed.

To check the status of the events you have registered:

1. Click the **Monitor** tab.

2. From the **Monitor** drop-down list, choose **Events**. Each registered event is listed in the *Monitor Events* page.

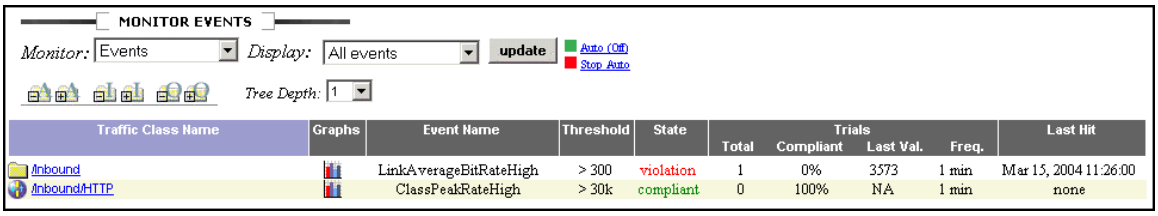

The columns on the *Monitor Events* screen are described in the following table.

Column	Description
Traffic Class Name	The full pathname of the class, link, or partition for which the event is registered
Graphs	The icon displays the event graph; see View Event Graphs.
Event Name	The event definition name
Threshold	A relative operator (>, >=,<, <=) followed by the threshold value that was defined when the event was registered
State	*compliant* displayed in green if the registration state is armed and the current value has not exceeded the threshold *disarmed* displayed in black if the event registration state is disarmed (an event is disarmed when the variable has exceeded the threshold value, but the event hasn't been re-armed yet; no emails or traps will be sent when an event is disarmed) *violation* displayed in red if the value exceeds the threshold
Trials Total	The number of times the event has been checked
Trials Compliant	The percentage of compliant trials, truncated to an integer
Trials Last Val	The value of the event definition variable the last time the event was checked
Trials Freq	The event-checking frequency

Column	Description
Last Hit	The last time the event was triggered, which is the last bad trial while the event was armed

View Event Graphs

In the *Monitor Events* window, each event has a link for displaying a graph. This graph shows the threshold alarm level and re-arm values as compared to the actual statistics of the event's measurement variable.

To view the graph for a registered event:

1. Click the **Monitor** tab.

2. From the **Monitor** drop-down list, choose **Events**. The *Monitor Events* page appears.

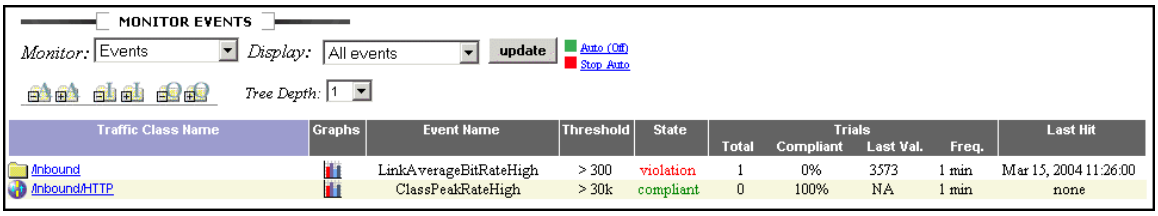

3. In the list of events, locate the event you want to graph.

4. For the event you want to graph, click the icon in the Graphs column.

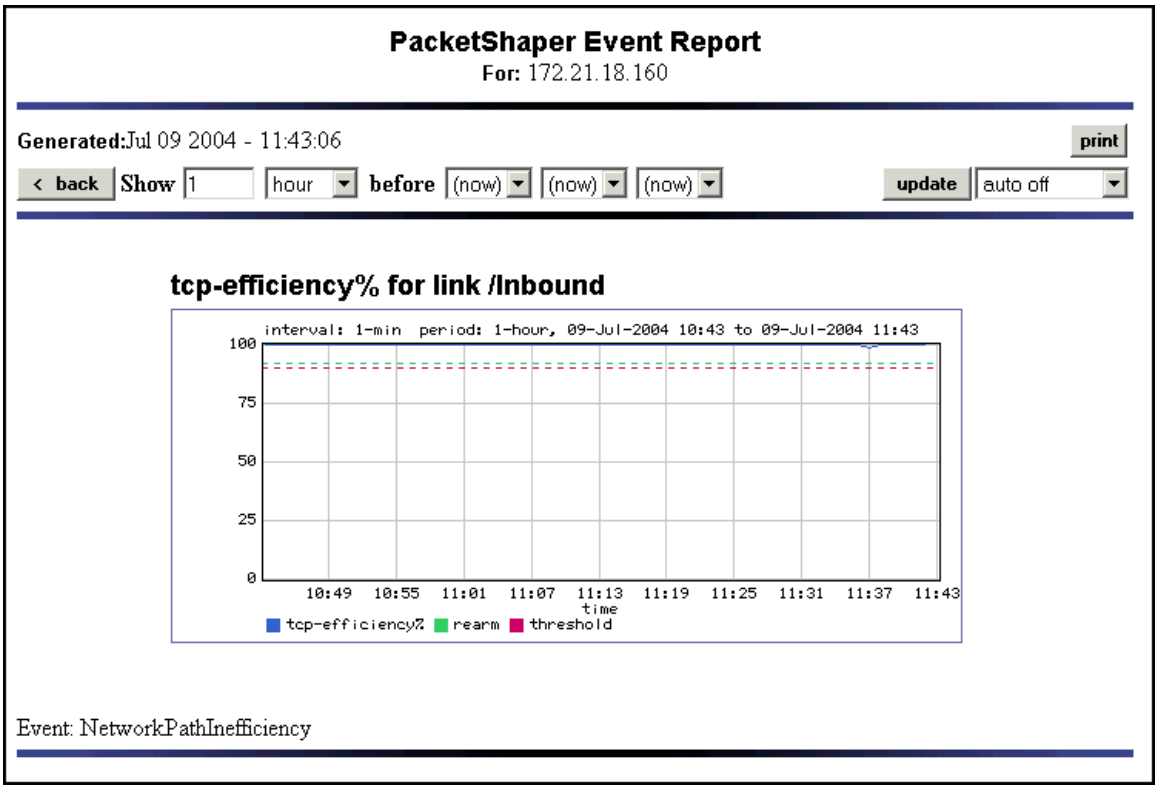

The event graph displays three colored lines that represent the following event statistics:

- Blue line – The recorded statistics for the event's measurement variable
- Red line – The threshold value for the event. The event is triggered when the threshold condition is met

- Green line – The re-arm value for the event. The event is re-armed after a threshold condition is met, so that another event can be triggered

View Event Logs

All event notifications are logged in files on the unit's data disk *before* each notification is sent. Event records are appended to the end of the current log file. The current log file is limited to a default size, so when it reaches capacity, old entries are archived.

To view the status of event-notification log files:

1. Access the CLI.

2. At the command-line prompt, type **event log status** and press Enter to display information about the current and archived event log files. (Output is shown below.)

Event log status: ok
Log file directory: 9.258/log
Maximum file size: 1152000 bytes (4000 event records)
Maximum number of archived log files: 4
Number of events in current log file 9.258/log/events: 0

To view the event-notification log files:

1. In the Advanced UI, click the **Info** tab and click **file browser**.

2. Go to the unit's data disk, 9.258/.

3. Go to the LOG/ directory. The event logs are listed. The file named EVENTS is the current set of log entries. The EVENTS.n files, where n is a sequential number, contain archived event records.

4. To view the contents of the current log file, click **EVENTS**.

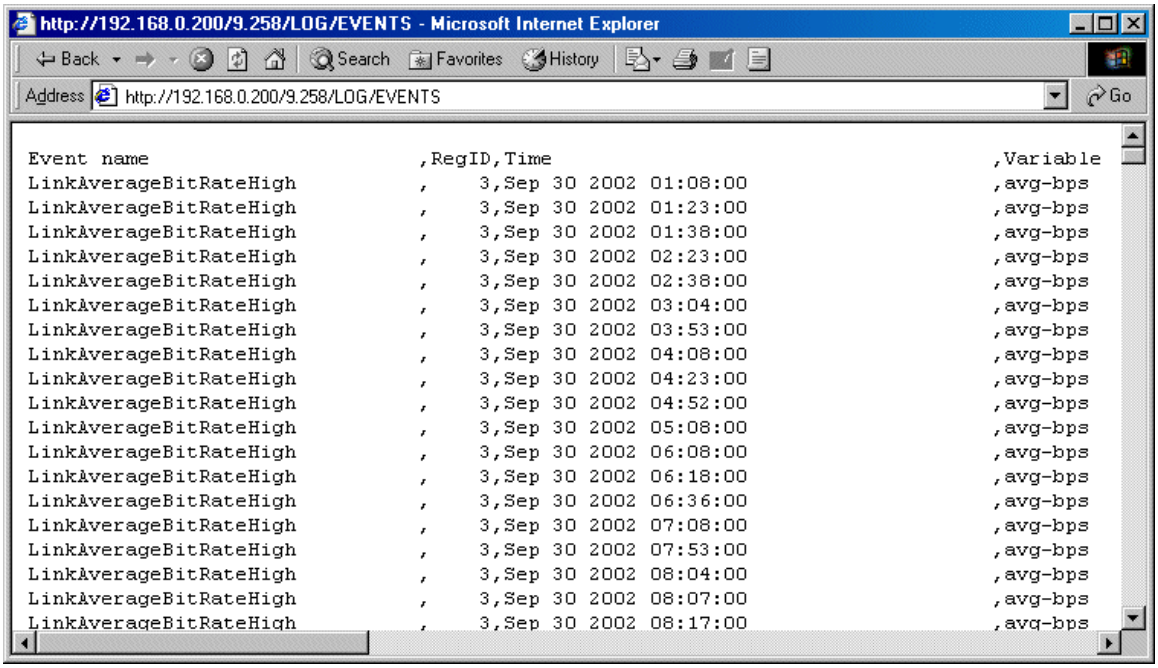

Each event log record has a 288-byte fixed length with comma-separated fixed-length fields. Each record is terminated with a carriage return character. The first record in the log file is a header record that identifies the field names and log-file version. The version is tracked so that in the future, if the record or file structure changes, PacketShaper can differentiate the records.

A log record lists the event name, ID, time, measurement variable, value, and object.

Manage the Email Queue

PacketShaper sends mail from its mail queue immediately. If the email server is down, PacketShaper retries at regular intervals. The retry facility tries to send the message three times, then doubles the interval and tries again three times. This process continues until it is successful or an hour has passed.

Use the email queue command from the CLI to list all messages that have not been sent yet, display the details for specific messages, or remove entries from the queue.

To bypass PacketShaper 's automatic retry facility and send messages immediately, use the email retry command.

Administration Tasks

PacketShaper offers several features that help you in the administration of a PacketShaper unit. There is a built-in file browser that allows you view, delete, zip, download, and upload files. In addition there is a command file editor and a facility for scheduling PacketShaper commands.

Use the File Browser

The PacketShaper File Browser is a convenient way to access files on the PacketShaper. With the File Browser, you can see what files and folders are on your unit, and view and edit certain types of files.

In addition, the File Browser has tools that allow you to:

- Delete files
- Create ZIP files
- Upload files to your PacketShaper
- Download files from PacketShaper
- Edit and create command files

To use the file browser:

1. In the Advanced UI, click the **Info** tab.

2. Click **file browser**. The File Browser window appears. **show screen...**

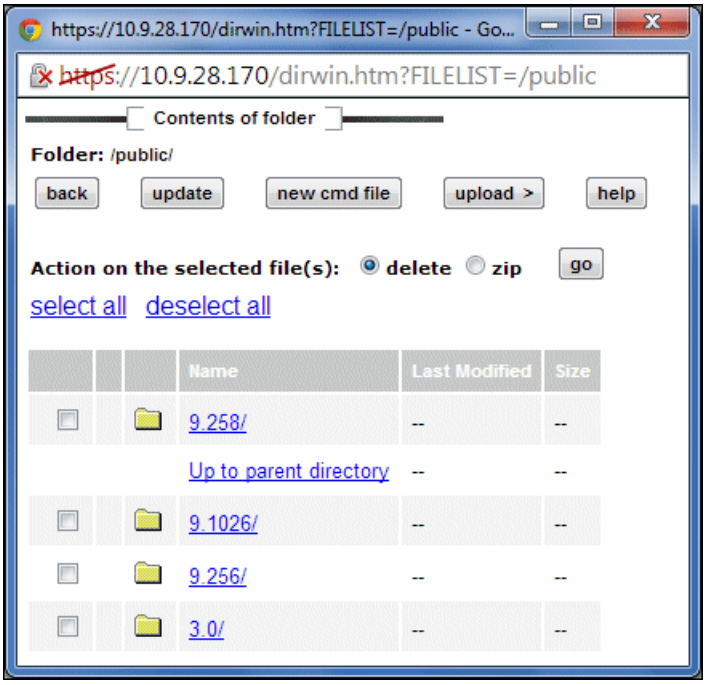

This screen lists the available public folders on the PacketShaper.

3. To change to another folder, click its name in the Name column. You may need to use the scroll bar to see additional folders. (See PacketShaper Folders for descriptions.)

Delete Files

With the File Browser, you can delete files from your PacketShaper. You can delete one or more files in the same folder in a single operation. You can delete folders, too, but only if they are empty.

To delete files on your PacketShaper:

1. In the File Browser, go to the folder where the file is located.

2. Select the checkbox next to each file you want to delete. **show screen...**

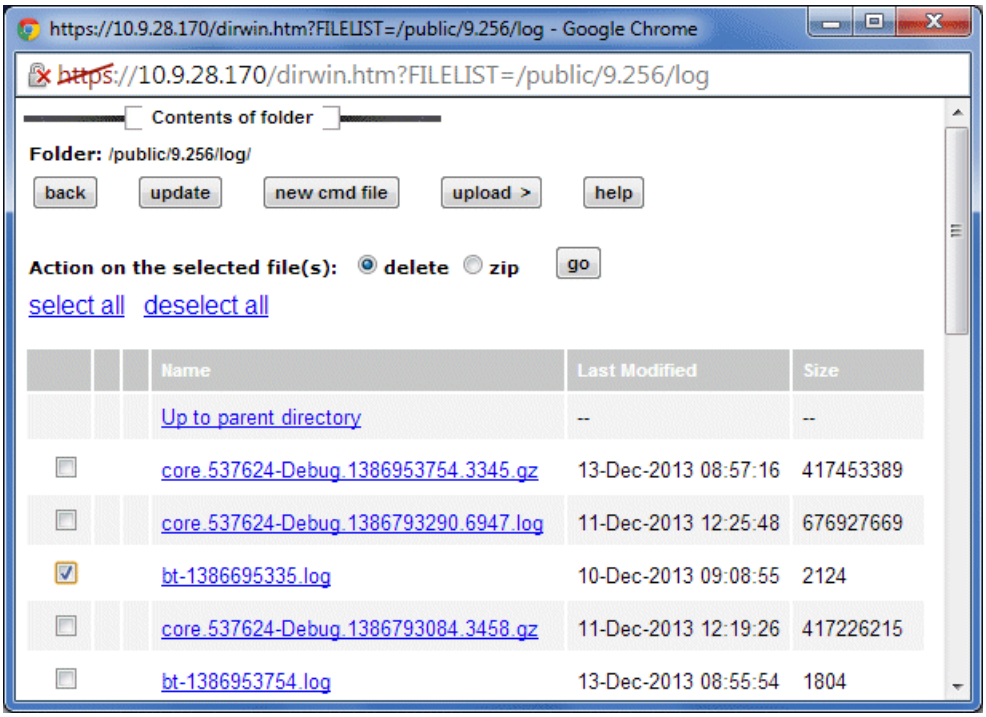

   or

   To select all files in the current folder, click **select all**.

3. For the action on the selected files, make sure the **delete** option is selected.

4. Click **go**. A status message, such as *Deleted TEST.TXT*, appears in the Results area at the bottom of the window.

5. If you still see the file listed in the *Contents* window, click **update** to refresh the screen.

Create a ZIP File

Using PacketShaper's File Browser, you can compress one or more files that are on your PacketShaper into a standard ZIP file before downloading them to a PC. For example, you can compress a set of measurement data dumps that have been exported to text files, or diagnostic logs in the 9.258/log folder. The PacketShaper zipping utility uses the recursive parameter: if a folder is selected to be zipped, the utility will automatically zip all of its contents, including any nested folders.

To zip files on your PacketShaper:

1. In the File Browser, go to the folder where the files are located.

2. Select the checkbox next to each file or folder you want to include in the ZIP file.

   or

   To select all files and folders in the current folder, click **select all**.

3. For the action on the selected files, select the **zip** option. **show screen...**

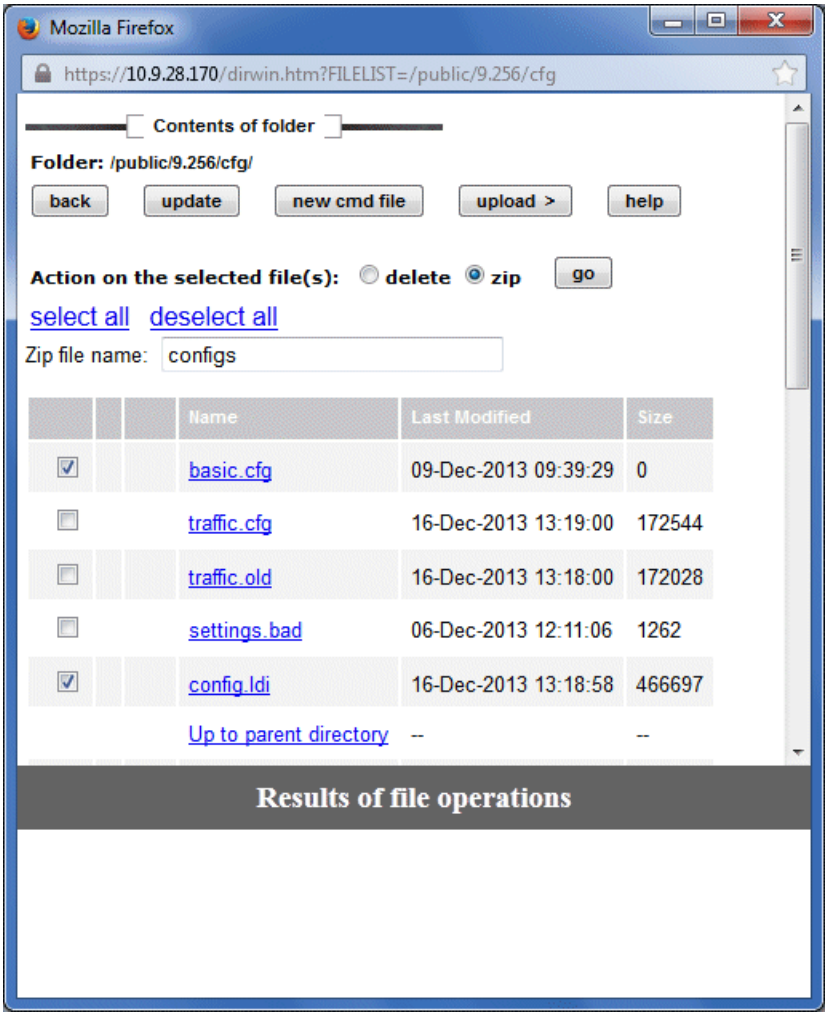

4. In the **Zip file name** field, type a descriptive name for the file.

   The .zip extension will be added automatically (for example, logfiles.zip). The file is stored in the current folder unless you specify a path.

5. Click **go**.

6. Read the summary message in the Results area at the bottom of the window. The message indicates which files were added and the percentage each file was compressed.

   The ZIP file will be listed in the Contents window. If you don't see the file, click **update** to refresh the screen.

Open a ZIP File

To view the contents of a ZIP file:

1. In the File Browser, click the name of the ZIP file.

2. When the download window appears, choose the option to **open** the file. **show screen...**

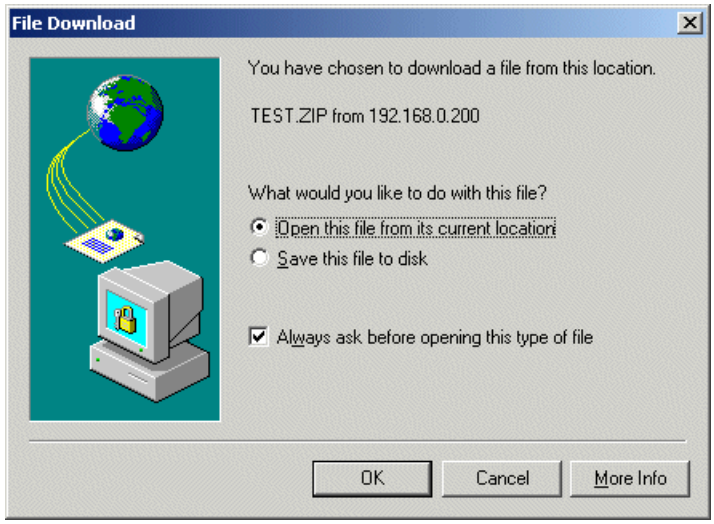

   The ZIP file is opened in your default ZIP utility.

3. You can now extract files, if you like.

4. Close the window when you're finished.

Unzip a File

When you unzip a file, all files in the ZIP file are extracted to the current folder. If you want to extract certain files only, you can open the file in your external ZIP utility (see Open a ZIP File).

To unzip a file:

1. In the File Browser, click the **unzip** button next to the name of the ZIP file. **show screen...**

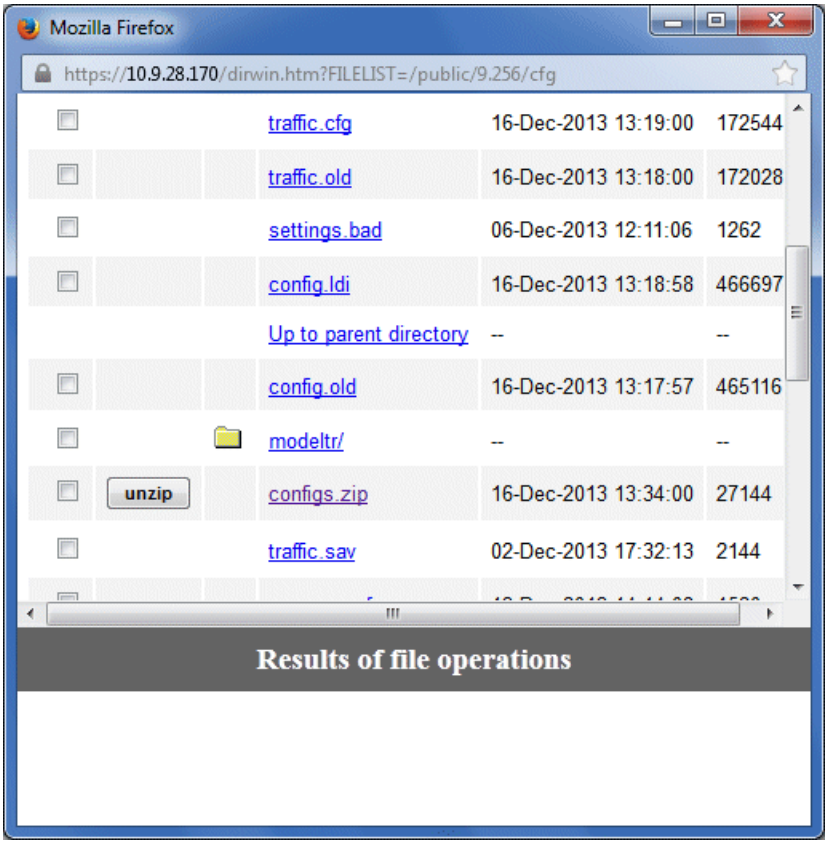

   A message appears warning you that any files contained in the ZIP file that already exist will be overwritten.

2. Click **OK** to confirm.

3. Read the message in the Results area at the bottom of the window to confirm that the files were extracted.

Download a File

You can use the File Browser to download (copy) a file from the PacketShaper to a PC or network drive.

To download a file from your PacketShaper:

1. In the File Browser, navigate to the folder containing the file to download.
2. Right-click the filename to download and choose the **Save as** option (such as **Save link as** in Chrome).
3. Save the file on your PC or network.

Upload (Copy) a File to PacketShaper

You can use the File Browser to upload a binary or text file from a PC or network drive to the PacketShaper.

To upload a file to your PacketShaper:

1. *(Optional)* In the File Browser, go to the destination folder.

   Although this step is not mandatory, doing this step will automatically fill in the destination when you give the upload command.

2. Click the **upload** button. The *File upload* screen appears. **show screen...**

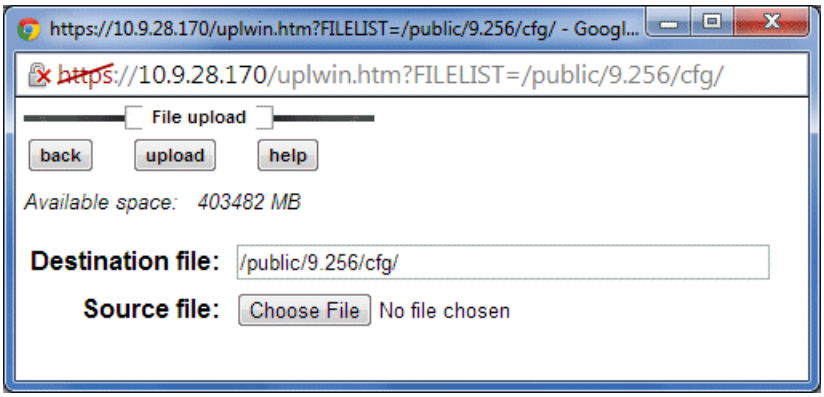

3. Check the available space on the current volume to make sure you have sufficient space for the file you are uploading.

4. Click **Browse** and select the file to upload.

   The **Source file** field is filled in with the name of the file you selected. In addition, the **Destination file** field is filled in with the current folder and the name of the file you are uploading.

5. If necessary, edit the path or filename in the **Destination file** field.

6. Click **upload**.

7. Read the summary message in the Results area at the bottom of the window and make sure the file was successfully uploaded.

8. To upload additional files, repeat steps 4 through 7.

9. Click **cancel** when finished.

   The uploaded files are listed in the Contents window. If you don't see the files, click **update** to refresh the list.

Note: When uploading a PacketShaper software image file, make sure to upload it to 9.258/image.

Create a Command File

When you have a sequence of commands that you type over and over again, you can save time by typing these commands once into a *command file*. A command file is a text file containing a list of PacketShaper CLI commands. When you run the command file, the commands are executed sequentially. You can use the run command to execute the command file from the CLI prompt, or use the schedule feature to schedule the command file to run on a regular basis.

> **Tip:** If you name the file startup.cmd and place it in 9.256/, the commands will be executed whenever the PacketShaper is rebooted or turned on.

Not all CLI commands can be included in a command file. First of all, there is a requirement that the CLI command run without additional user input. If the command is interactive and prompts the user for input, the command won't work. Second, many of the file utility and diagnostic commands (such as **cat**, **cd**, **cmp**, **date**, **dns lookup**, **du**, **echo**, **ls**, **mkdir**, **mv**, **ping**, and **rm**) cannot be included in a command file. Exceptions are **zip** and **unzip** these commands can be included and are quite useful in command files.

Although you can create a command file in any text editor, you may prefer the convenience of using PacketShaper's Command File Editor. With the built-in editor, you can save the file directly on the PacketShaper unit.

To create a command file:

1. In the File Browser, navigate to the folder where you want to store your command file.

2. Click the **new cmd file** button. The *Command File Editor* window appears. **show screen...**

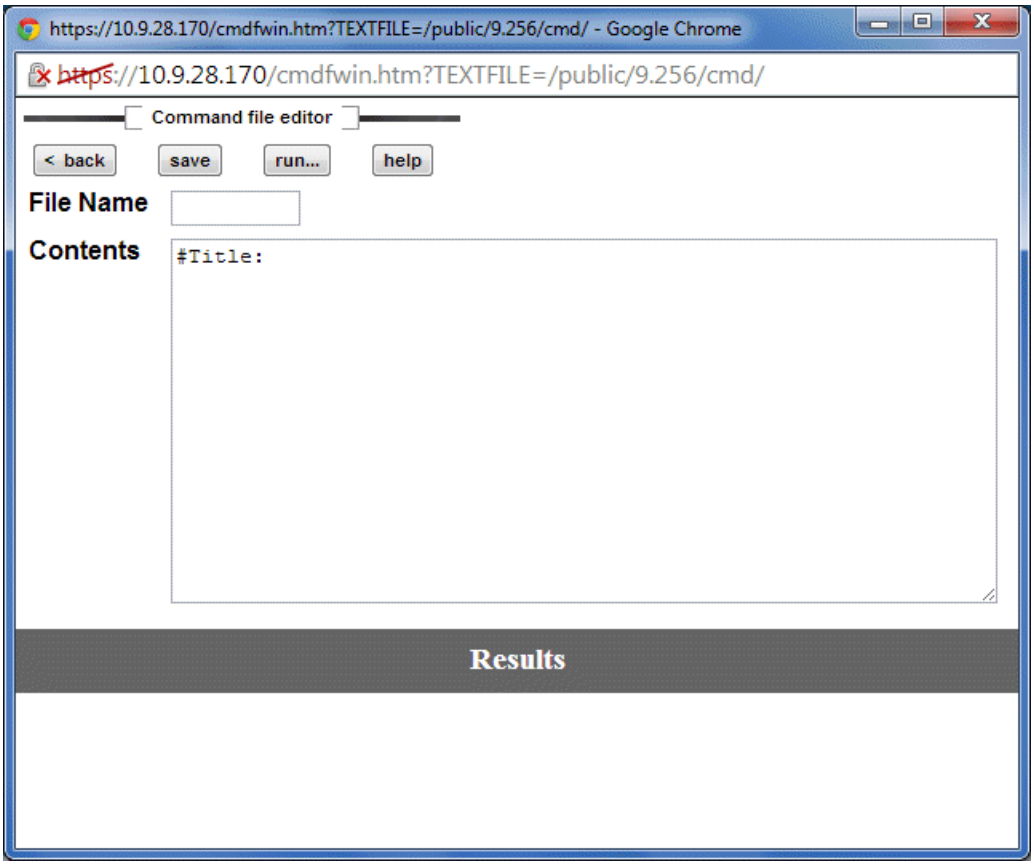

3. *(Optional)* In the Contents area, enter a description next to the *#Title* text. You can use this comment line to describe the function of the command file.

4. Enter the CLI commands, with each command on a separate line. (See Command-Line Interface for descriptions and syntax of each command.)

> **Note: Important:** Command files must end with a carriage return. If a command file doesn't have a carriage return at the end of the file, the last command in the file will not be executed. Therefore, make sure all of your command files have a carriage return after the last line.

5. To save the file, enter a descriptive name (up to 8 characters) in the **File Name** field and click **save**.

The file is stored in the current folder and a CMD file extension is automatically added to the name. The Results area indicates whether the file was saved successfully.

6. To make sure you entered the commands with the correct syntax and that the commands work as expected, click **run**. The Results area displays the output of each command and any error messages that may have occurred.

7. Modify the commands as necessary, and then repeat steps 5 and 6 until the command file runs successfully.

8. Click the **back** button and observe that the command file you created is listed in the *Contents of folder* window. If it's not listed, click **update** to refresh the screen.

If you need to modify the command file in the future, use the **edit** button that appears next to the filename. Note that the maximum file size that can be edited is 8K.

Review Log Files

Each time you run a command file, a log file is created of the command output. The log file contains the same text you see in the Results area of the *Command File Editor* window after you run the command file.

To view the log files:

1. In the File Browser, go to the folder where the command files are located.

2. Open the COMPLETE folder.

This folder contains all the log files for executed command files. A log file has the same name as the command file, but with a numbered extension. For example, a command file named CUST.CMD would produce log files named CUST.1, CUST.2, and so forth. You can use the date and time in the Last Modified column to help you figure out which log file you want to open.

3. To view a log file, click its name. The text file appears in the same window.

4. Click the browser's **Back** button to return to the *Contents of folder* screen.

It is recommended that you periodically go to the COMPLETE folder and delete old log files. (They are not deleted automatically.) See Delete Files.

Create a Schedule

With the scheduling feature, you can schedule commands to be executed at a specific time, on a specific day, on an ongoing basis or only once. This feature issues PacketShaper commands when your PacketShaper is unattended, allowing you to automatically vary bandwidth-allocation policies according to the time of day or day of the week. For example, you can schedule a policy that restricts MP3 file downloads during business hours and schedule another less restrictive policy for after-hour music downloads. This feature could also be used to schedule tasks (such as backing up measurement data) to run during a specified maintenance window. You can create up to 64 schedules.

You can schedule a single CLI command or a series of commands. If you want to create a scheduled item for multiple commands, you need to enter the CLI commands into a text file. For more information, see Create a Command File.

To create a new schedule:

1. Click the **Setup** tab on the navigation bar.

2. From the **Choose Setup Page** list, select **command scheduling**.

3. Click **new**. The *Add Schedule Entry* page appears.

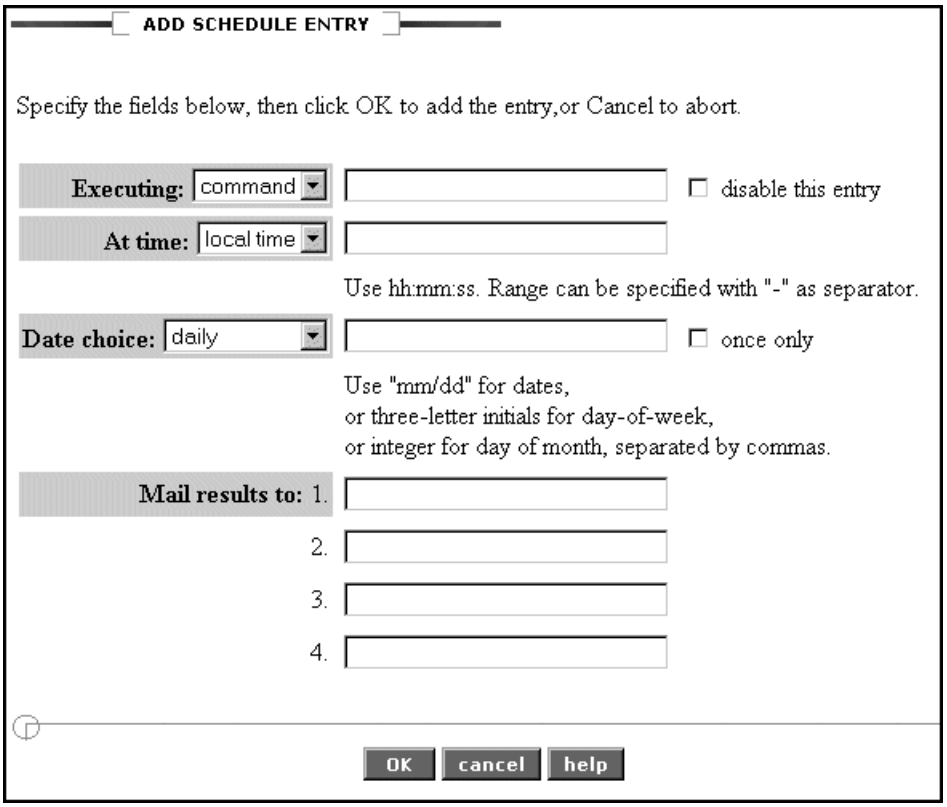

4. In the **Executing** field, specify the command(s) you want to schedule:

Choose **command** and enter the CLI command in the empty text field. (See Command-Line Interface.)

-or-

Choose **file** and enter the name of the command file in the empty text field. If the file is not located in the 9.256/cmd folder, specify the complete path.

6. In the **At time** field, specify the time you want to execute the command. Select one of the following:

Choose **local time** and enter the time in the format *hh:mm:ss* where *hh* is the hour from 0 to 23, and *mm* and *Ss* are minutes and seconds from 0 to 59. For example, to specify the time 5pm, enter 17:00.

-or-

Choose **UTC time** and enter a coordinated universal time (UTC). Specifying times in UTC (a standard similar to Greenwich mean time) is useful when managing units in different time zones. For example, 1800Z is 1pm in Eastern standard time and 4pm in Pacific standard time.

> **Note:** PacketShaper will attempt to execute the command within a 40-second window. If you want to allow more time for the command(s) to be executed, you can specify a range (for example, 08:00-08:02).

> **Tip:** When setting the time for a scheduled command, keep biannual time changes into consideration. For example, if you set a command to execute at 2:30am in the United States, the command will not be executed when the clock changes ahead one hour for Daylight Saving Time. You can ensure a command will be executed during a clock change by specifying a time range (such as 02:00-04:00).

7. Select the day(s) you want to execute the command:

Option	Description (when the command is executed)	Additional Specification (entered in the adjacent text field)
daily	Every day of the week	n/a (no further specification is needed)
dates	On a specific date or series of dates	Specify the date in the format *mm/dd*. The date is assumed to be a future date, within the next twelve months. For example, if today is 5/30/02 and you specify the date 5/29, the item will be scheduled for execution on May 29, 2003. You can specify up to 10 dates separated by commas.
weekday	On weekdays only (Monday through Friday)	n/a

Option	Description (when the command is executed)	Additional Specification (entered in the adjacent text field)
weekend	On weekends only (Saturday and Sunday)	n/a
day-of-week	On specific days of the week	Specify the day using the first three letters of the day (mon, tue, wed, thu, fri, sat, sun). If you specify more than one day, each day must be separated by a comma and no space, for example, *mon,wed*. You can specify up to seven days of the week.
day-of-month	On specific days of the month	Specify the integer day of the month for example, *15* for the fifteenth of the month. You can specify up to 31 days, separated by commas.
today	On today's date	n/a

8. If you want the command to be executed once for each of the specified dates, select the **once only** check box. If you don't select this check box, the command will be executed on an ongoing basis, according to the date(s) you specified.

9. In the **Mail results to** text fields, specify up to four email addresses to which you want to send the output of the command or command file execution. This allows you to confirm that the command executed at the specified date and time.

> **Note:** In order to send email confirmations, you must configure an SMTP server.

9. Click **OK**. The scheduled item is now listed on the Command Scheduling setup page.

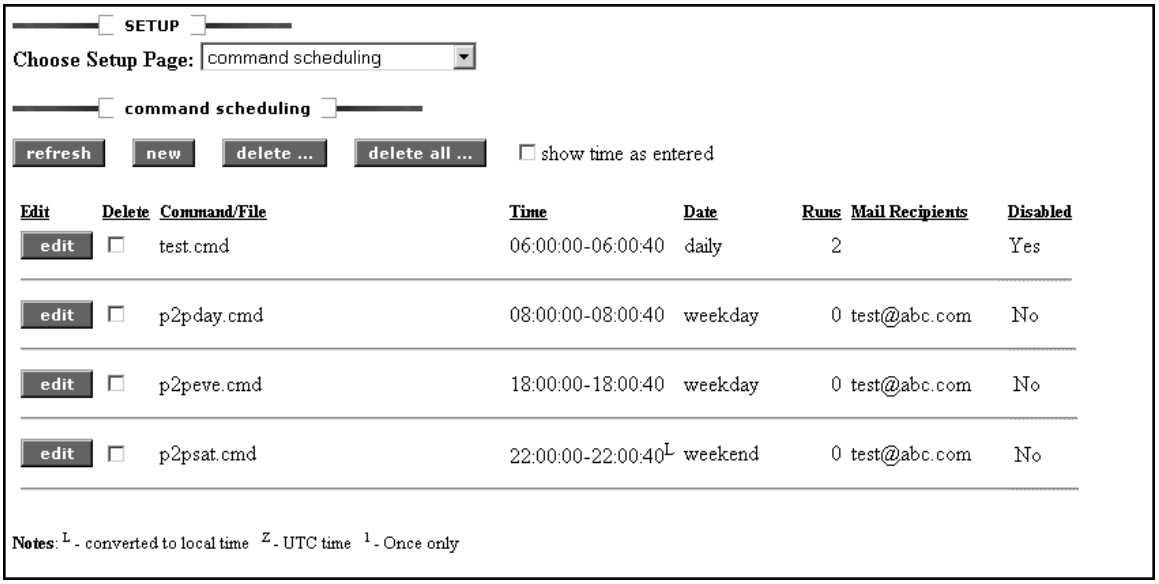

Scheduled commands are part of the unit's configuration and are retained after resetting or rebooting the unit.

List Schedules

To list all scheduled items:

1. Click the **Setup** tab on the navigation bar.

2. From the **Choose Setup Page** list, select **command scheduling**. The list of scheduled items appears.

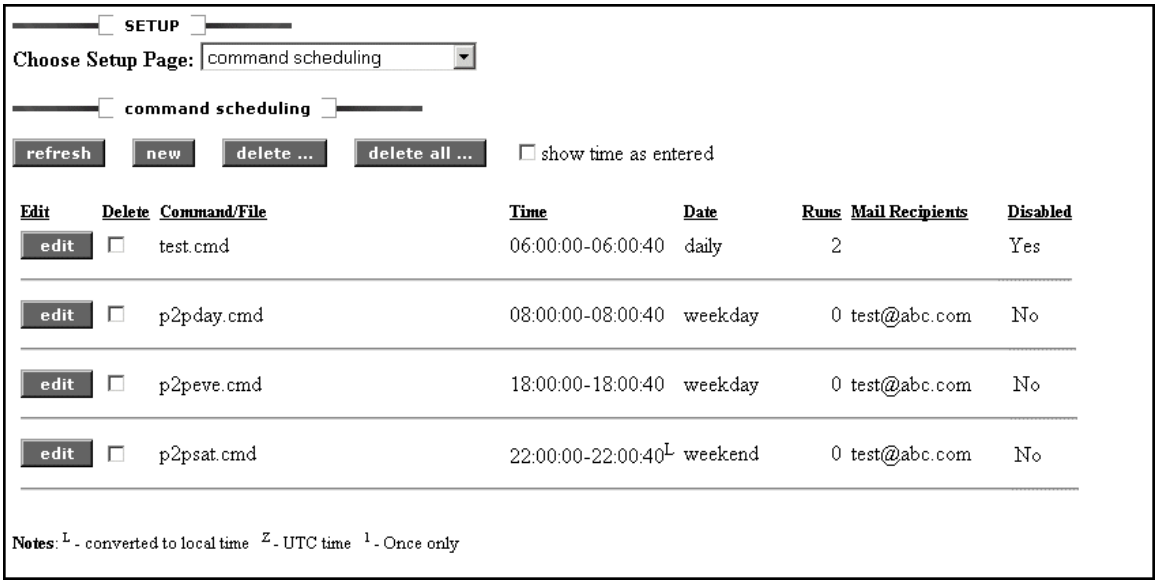

The Command Scheduling setup page lists the commands for which you have created a schedule and provides buttons for editing and deleting the scheduled items. The table has the following columns:

Column	Description
Edit	Modifies the scheduled item. See Edit a Schedule.
Delete	Removes the scheduled item. See Delete a Schedule. **Note:** Items that are scheduled to run only once are automatically removed after they are executed.
Command/File	The command or command file to be executed
Time	The scheduled time that the commands will be executed. All times are listed in local time. If the scheduled time was entered in UTC (coordinated universal time) format, it is converted to local time on the display and an **L** appears next to the time, indicating that the time is shown in local time format. **Note:** If you want to list the times as they were originally entered – in UTC and/or local time select the **show time as entered** check box. Any times that were entered in UTC format will have a **Z** next to them.
Date	The day(s) the item is scheduled to execute. If the item is scheduled to execute only once, a **1** appears next to the date.
Runs	The number of times the schedule has run
Disabled	Indicates whether the item has been disabled. (To disable an item from executing, edit the schedule and select the **disable this entry** check box.)

Column	Description
Mail Recipients	The email addresses to which a confirmation message will be sent after a command has executed.

Delete a Schedule

One-time-only schedules, such as those with the "today" date option, are automatically removed after they are executed. Scheduled items that are ongoing, however, must be removed manually if you want to cancel them.

If you think you might want to use a scheduled item again, you can disable it instead of permanently deleting it. See Edit a Schedule for details on disabling a scheduled item.

Note: To permanently delete a command scheduled via PolicyCenter, you must remove the command from the unit's sharable PolicyCenter configuration, and not just from the unit itself. If PolicyCenter detects that a scheduled command is on the unit's PolicyCenter configuration but no longer on the unit itself, PolicyCenter will synchronize the unit's settings with its PolicyCenter configuration to restore that command to the unit. If the time range for the scheduled command is not over or the PolicyCenter time zone is not correctly configured, the command may run again.

To delete a schedule:

1. Click the **Setup** tab on the navigation bar.

2. From the **Choose Setup Page** list, select **command scheduling**.

3. For each item you want to remove, select the check box in the Delete column. **show screen...**

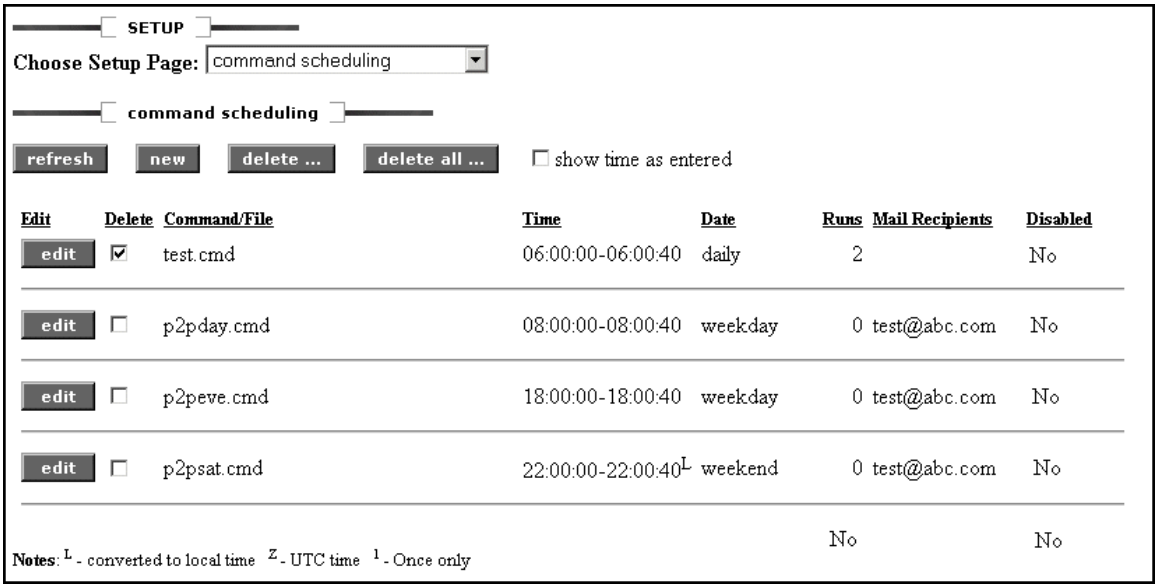

4. When you're finished marking the items to remove, click the **delete** button.

To delete all scheduled items, you can simply click the **delete all** button – you don't need to mark each item.

Edit a Schedule

You can modify any aspect of a scheduled item: the command, time, and date.

To modify a schedule:

1. Click the **Setup** tab on the navigation bar.

2. From the **Choose Setup Page** list, select **command scheduling**.

3. For the item you want to modify, click **edit**. The *Edit Schedule Entry* page appears.

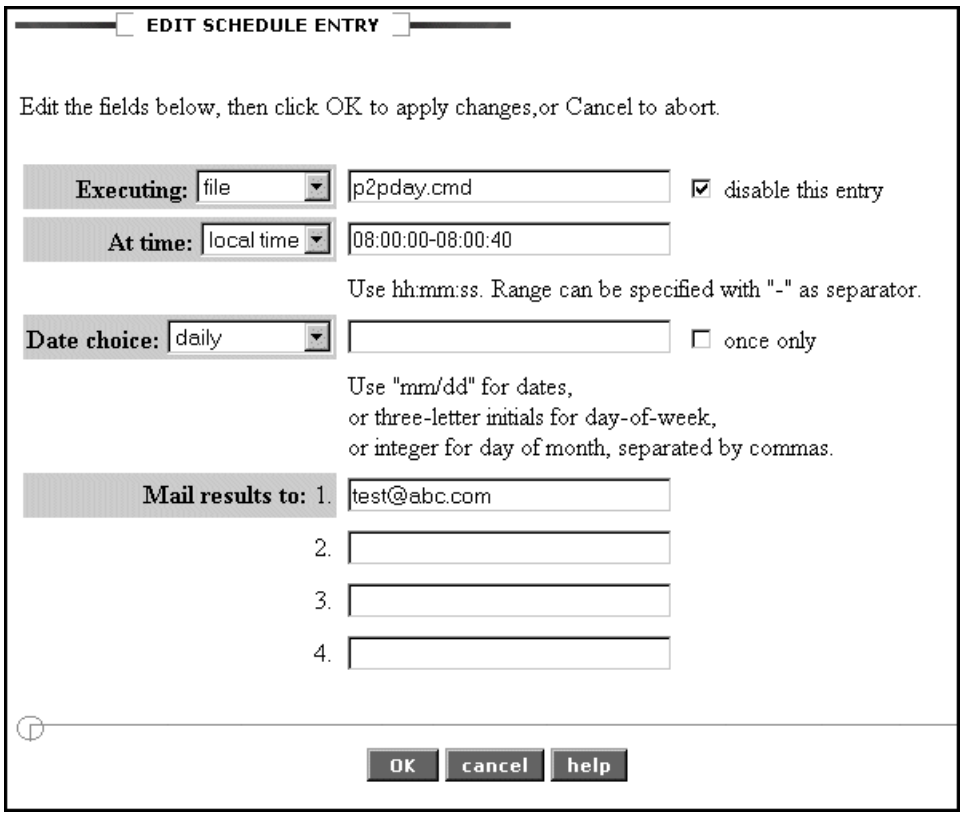

4. Modify the fields as necessary. For detailed descriptions of the fields, see Create a Schedule.

5. To disable the schedule so that it won't execute, select the **disable this entry** check box.

> **Note: Note:** If you want to enable it later, unselect this check box.)

6. Click **OK**. The *Command Scheduling* page appears and displays the new settings for the scheduled item you modified.

Back Up a Configuration

There are a number of reasons for backing up your standalone (local mode) PacketShaper configuration:

- Your configuration may not load properly after upgrading to a new version of PacketShaper software.
- You could upgrade to a new release and then realize that you want to revert to a previous version.
- You might need to exchange your current PacketShaper for a different unit (for example, a newer model) and need to migrate files to the new unit.
- You can preserve the current configuration and use it to configure other units.
- You will be able to recover critical configuration files in the unlikely event of file corruption.

PacketShaper automatically stores your settings in a file named config.ldi. In addition, you can save a configuration on demand, using the **config save** CLI command. This command creates two files. The .ldi file contains the traffic tree configuration (including all classes, class IDs, partitions, policies, host lists), as well as all sharable configuration settings such as packet shaping, traffic discovery, passwords, SNMP, SNTP, and Syslog. The .cmd file contains the nonshareable settings such as IP address, gateway, date, time, and network interface settings.

You can also capture your traffic configuration and settings in an executable command (.CMD) file, using the **setup capture** CLI command.

Note that restoring a configuration by running a CMD file takes much longer (possibly hours) than loading a configuration (less than a minute). However, Symantec recommends that you create and back up the CMD file as a safeguard in case the config.ldi fails to load.

Capture the Traffic Configuration

The **setup capture** command saves your traffic configuration and settings in an executable CMD file. If you ever need to restore the settings you captured, you can use the **run** command to re-create the configuration.

To save your settings in a CMD file:

1. Access the command-line interface.
2. Decide whether to capture unit-specific details, such as the setup commands for the IP addresses. The *complete* format stores unit-specific settings; the *portable* format does not.
3. Enter one of the following commands:

   ```
   setup capture portable <filename>

   setup capture complete <filename>
   ```

   where *<filename>* is the name of the CMD file (such as *backup.cmd*); it is not necessary to type the .CMD file extension.

This ASCII file is saved in the 9.256/cmd folder and should be backed up to a client workstation.

Save the Configuration

Save the current configuration's sharable settings in an .ldi file and its nonsharable settings in a .cmd file. Shareable settings include the traffic tree configuration (all classes, class IDs, partitions, policies, host lists), as well as all sharable configuration settings such as packet shaping, traffic discovery, passwords, SNMP, SNTP, BCAAA, adaptive response, and Syslog settings. Nonshareable settings include IP address, gateway, date, time, and network interface settings.

A single CLI command creates the .ldi and .cmd file. Enter the following command:

`config save <filename>`

where *<filename>* is the name of the configuration file. The .ldi and .cmd extensions are automatically added to the configuration file name.

Back Up the Configuration Files

After creating the configuration backup files, you should copy them to a workstation's hard drive.

To transfer files from the PacketShaper to a workstation hard drive:

1. At your workstation's command line (e.g., MS-DOS prompt), create a folder where the backup files will be stored. Make sure that the folder name identifies the specific unit so that the .ldi file stored therein can be used to recover the unit.
2. Open the PacketShaper's File Browser.
3. Go to the 9.256/ folder and download the .ldi and .cmd text files you created with the **config save** command.
4. Go to the 9.256/CMD folder and download the CMD file you created with the **setup capture** command.
5. *(Optional)* Go to the 9.256/CFG folder and download the conflig.ldi text file. This file can be backed up as an extra precaution, and could be useful when the unit is connected to PolicyCenter.
6. Confirm the backup files are stored in the folder you created in step 1.

Restore a Configuration

In the event that your current PacketShaper configuration becomes corrupt or if you want to revert to a previously saved configuration, use the following procedure to restore a configuration backup:

1. Open the PacketShaper's File Browser.

2. Go to the 9.256/ folder and upload the .ldi and .cmd files (created with the **config save** command) from the backup on your workstation.

3. To load the new configuration, go to the PacketShaper s command-line interface, and type the following command:

   `config load <filename> complete`

   or

   `config load <filename>` if you do not want to load nonsharable/unit-specific settings

   where *<filename>* is the name of the configuration file. You do not need to include the file extension.

4. If a configuration won't load or the traffic tree still isn't in place, you can restore the configuration by running the CMD file you backed up. For example, if you used the **setup capture** command and created a file named backup.cmd, you need to upload the backup.cmd file to the PacketShaper and then type **run backup.cmd** at the CLI prompt. To look for any errors, type **more backup.out**.

Watch Mode Tasks

When a PacketShaper is in watch mode, it passively monitors the traffic on your network and performs all the traffic classification and reporting tasks that PacketShapers typically do. Since its role is a passive watcher, the PacketShaper does not perform any traffic shaping functions. Watch mode is used when you don t want the PacketShaper to be cabled into the main data path when your situation requires that the unit monitor the traffic non-inline.

For more details on the watch mode feature, see Watch Mode Overview.

Configure Watch Mode

To configure watch mode:

1. Select the appropriate non-inline topology.

2. Connect a cable to the SYS MGMT port.

> **Caution:** Do not connect both cables until after watch mode is enabled (step 4 below). If the PacketShaper is configured inline when you enable watch mode, network connectivity problems will occur. This is because in watch mode packets are dropped and not forwarded through the network.

3. If your PacketShaper hasn't been configured yet, run Guided Setup. To access the PacketShaper, use a workstation that is on the network connected to the cabled port. Alternatively, you can connect directly to the unit using the CONSOLE port.

4. Enable watch mode:

   - Click the **Setup** tab.

   - Set **Shaping** to **watch**. **show screen...**

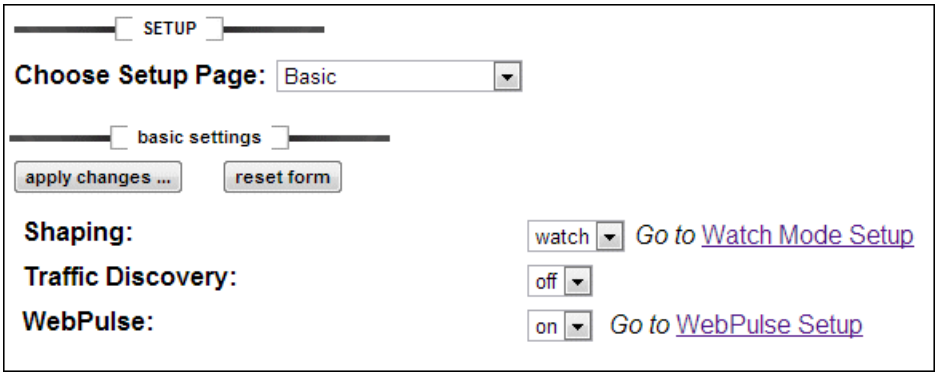

   - Click **apply changes**.

     After watch mode is enabled, the Info page says "*Shaping: Watch*."

5. If applicable to your topology, connect the other cable(s) to PacketShaper.

6. Define watch mode routers.

Select a Non-Inline Topology

In watch mode, PacketShaper will be connected to a switch's SPAN (Switched Port Analyzer) port so that it can monitor the WAN traffic. To monitor traffic sent by the switch to the router, the bidirectional traffic on the switch's router port is copied to the SPAN port so that it can be received and monitored by the PacketShaper.

> **Caution:** Do not actually make the connections shown in the following topologies until you have read all the steps in Configure Watch Mode.

PacketShaper can monitor a single LAN or, if additional NIC option cards are installed, it can monitor multiple LAN segments.

Select the topology that corresponds to your network. You'll need to refer to the diagram when connecting cables.

- Example 1: Monitor One LAN
- Example 2: Monitor Multiple LANs

After you have selected the appropriate topology, return to the Configure Watch Mode task.

Define Watch Mode Routers

The watch mode feature can monitor traffic to/from up to 256 routers. The PacketShaper uses the router list to determine direction of the packets the unit is monitoring: if a packet is *from* the router, it will be considered to be inbound or if the packet is *to* the router, the packet will be considered outbound. If both the source and destination addresses of a packet are in the watch list, the packet is considered same side traffic.

To add a router to the watch list:

1. Click the **Setup** tab and choose **watch mode** from the **Choose Setup Page** drop-down list.

2. Click **add router**. The *Add Watch Mode Router* screen appears. **show screen...**

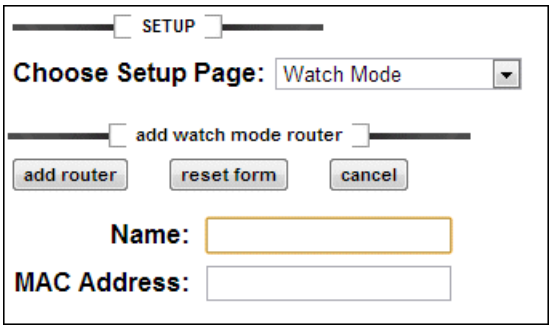

3. In the **Name** field, enter a descriptive name to uniquely identify the router. Names can be up to 32 alphanumeric characters and can contain certain special characters (colon, dash, underline, and period). Spaces are not allowed.

4. In the **MAC Address** field, enter the router s MAC address.

5. Click **add router**. The new router will appear in the *Watch Mode Router* list. **show screen...**

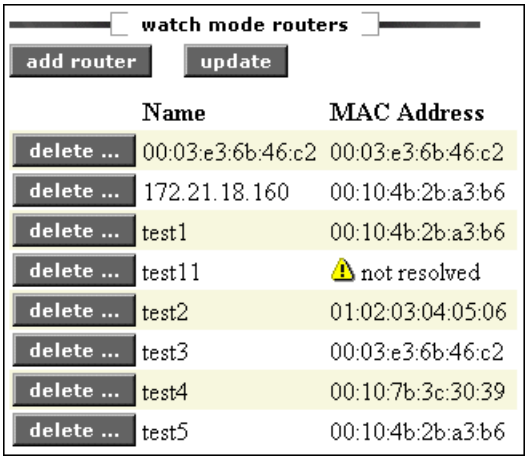

6. Repeat steps 2-5 for each router you want to add to the watch list.

Display Watch Mode Configuration

To view your watch mode settings, such as the list of routers configured for watch mode and the current management port, you can display the *Watch Mode Configuration* screen. Settings can be changed on this screen, as well.

To display the watch mode configuration:

1. If you are on the Basic Settings setup page, click the **Watch Mode Setup** link.

   or

   Click the **Setup** tab and choose **watch mode** from the **Choose Setup Page** drop-down list. The *Watch Mode Configuration* screen appears. **show screen...**

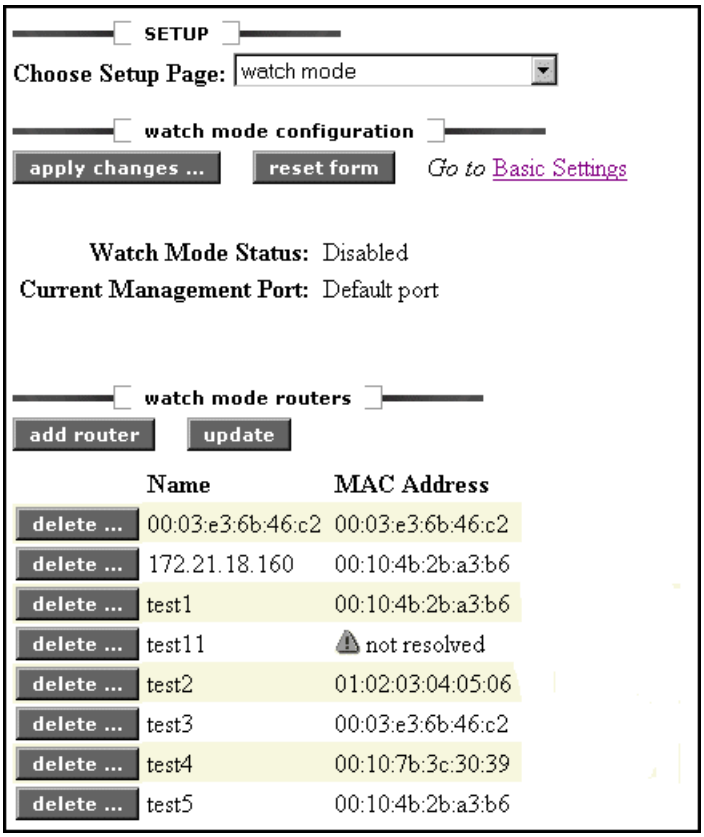

2. The following information is displayed in this screen:

Field	Description
Watch Mode Status	Indicates whether watch mode is enabled or disabled; to change the status, see Configure Watch Mode or Disable Watch Mode
Current Management Port	Indicates which port PacketShaper has determined will be used to manage the unit. The management port is not user-definable. PacketShaper decides which port to use for management access by checking which ports are connected.

Field	Description
Watch Mode Routers	The list of configured routers: Name and MAC Address. Use the **add router** button to configure a router for watch mode. Use the **delete** button to remove a router from the watch mode configuration. Use the **update** button to refresh the screen. **Note:** You cannot edit a router's settings; you must delete it and then add it again.

Delete Watch Mode Routers

To remove a router from the watch list:

1. Click the **Setup** tab.

2. Choose **watch mode** from the **Choose Setup Page** drop-down list. The currently configured watch mode routers are listed. **show screen...**

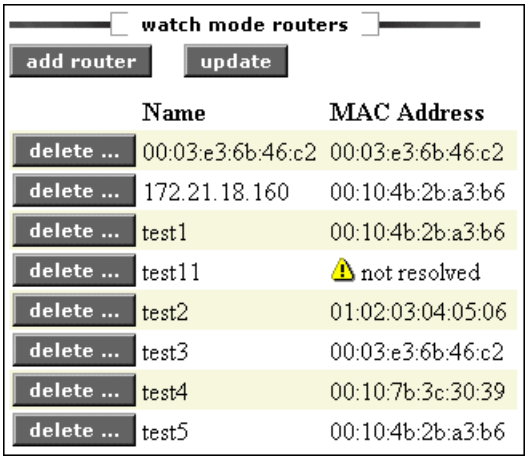

3. Next to the router you want to remove, click **delete**.

4. Confirm the deletion. The router is removed from the list.

Disable Watch Mode

To disable watch mode, you select **off** for the shaping mode. Any defined watch mode routers will remain configured, so that you don't need to redefine them if you later want to re-enable watch mode.

To turn off watch mode:

1. Disconnect the cable between the PacketShaper and the SPAN port–be sure you can access the PacketShaper through the SYS MGMT port.

2. Click the **Setup** tab. The Basic Settings screen automatically appears in the *Setup* window. **show screen...**

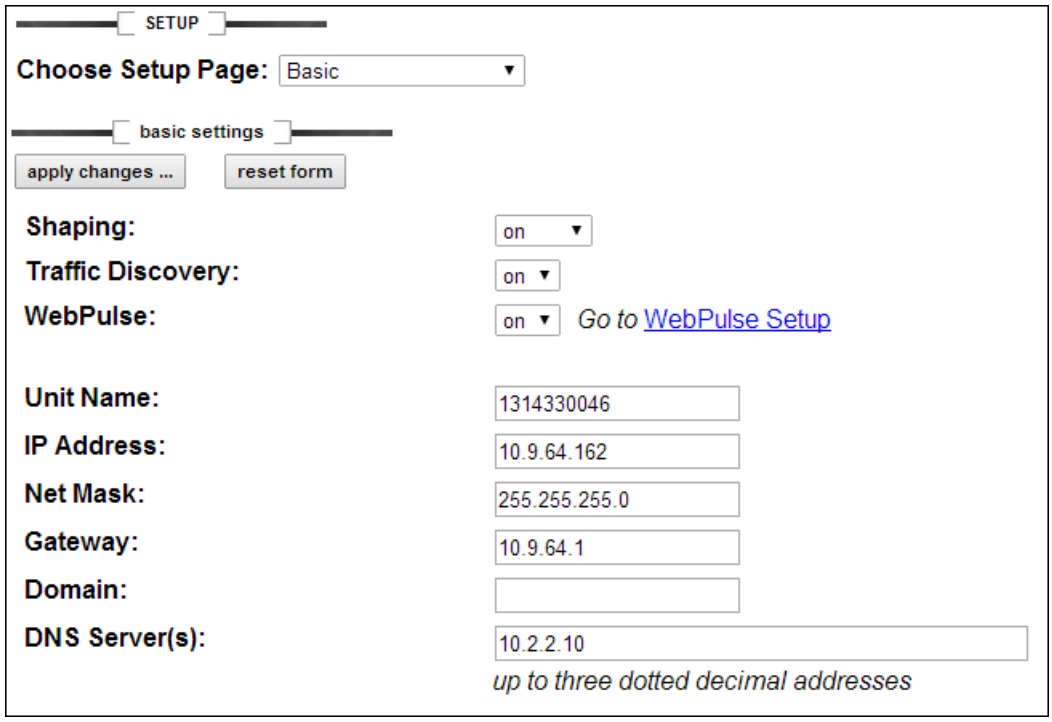

3. Set **Shaping** from **watch** to **off**.

4. Click **apply changes**.

After you disable watch mode, the Info page no longer says *Shaping: Watch*.

Connect to a Network Tap

Another way to place a PacketShaper non-line is to connect it to the network through a tap. A network tap is a transmit-only hardware device, placed inline, that provides a permanent access port for passive network monitoring. When a PacketShaper is connected to a tap, it receives the same traffic as if it were located directly on the wire.

The PacketShaper's INSIDE and OUTSIDE ports are connected to the tap – and **watch mode is not enabled**. The tap will copy traffic in both inbound and outbound directions.

Since the tap does not pass traffic back to the network, you must use the SYS MGMT port to manage the unit.

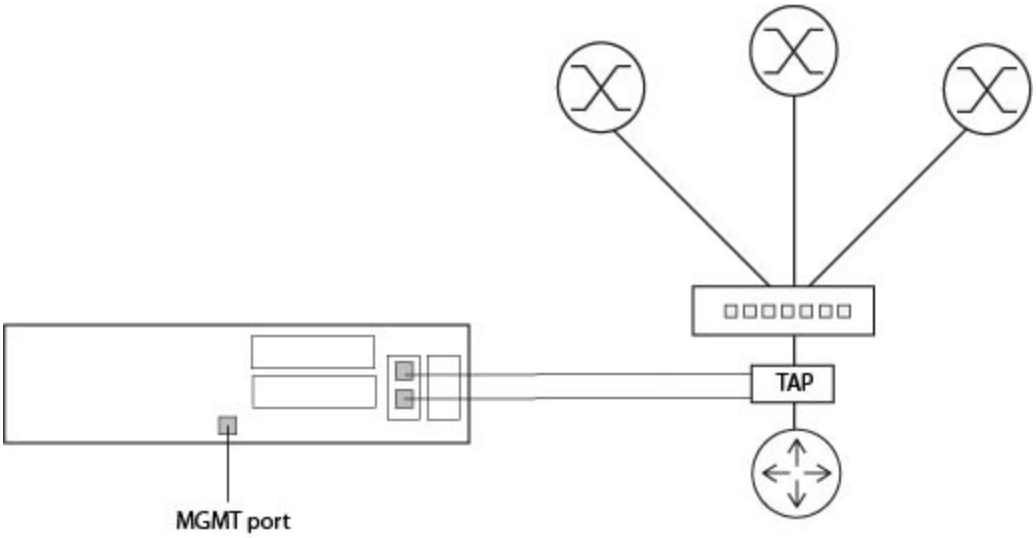

High Availability

PacketShaper offers two features for ensuring high availability:

- Standby
- Access-link monitoring

Configure Standby

The *standby* function allows two to five PacketShapers to work together in a redundant network topology. All partners in a high availability cluster are considered active and each unit can receive and forward traffic. As a result, each unit is on standby, ready at any time to take over another PacketShaper's responsibility should the unit go down.

1. Configure redundancy protocol on the switches and routers that are attached to the PacketShapers in the cluster.
2. Before enabling standby, make sure you have *not* set the outside interface to **secure**. For standby to work, each device must be able to communicate with the other device. If you have set the outside interface to **list**, you must add the partner's IP address to the Outside security list.
3. For each PacketShaper in the cluster, enter the following CLI command to set bypass to open: **setup bypass open all**. When bypass is open, the PacketShaper becomes a point of failure in the traffic path; this setting is required when using the standby feature.
4. Confirm that each appliance can ping the other PacketShapers that will be in the cluster.
5. On each PacketShaper, configure the IP addresses of its standby partners. This group of PacketShaper nodes becomes a high availability cluster. See " Configure PacketShapers for Standby" on the facing page.
6. Verify standby partners are connected. See " Verify Standby Partners" on page 724.

Configure PacketShapers for Standby

On each PacketShaper, use the Standby setup page to enable the feature and specify the IP addresses of the other PacketShaper nodes. This group of PacketShaper nodes becomes a high availability cluster.

1. Log in to one of the PacketShapers in the high availability cluster.
2. Go to the Advanced UI.
3. Click the **Setup** tab.
4. From the **Choose Setup Page** list, select **Standby**. The Standby settings are displayed.

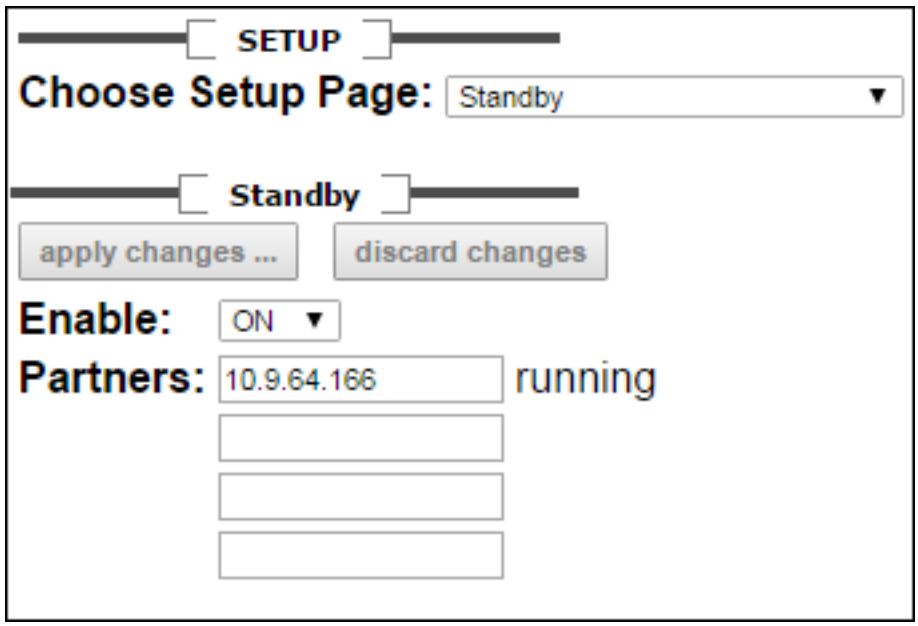

5. In the **Partners** fields, enter the IP address of each partner in the cluster.
6. For **Enable**, select **ON**.
7. Click **apply changes**.
8. Repeat the above steps for each PacketShaper in the cluster.
9. Verify standby partners are connected. See " Verify Standby Partners" on the next page.

Verify Standby Partners

After configuring standby partners and enabling Standby, you should verify that the partners are connected to one another and that you can view partner reports and see updates made on the partner.

1. On a partner in the high availability cluster, go to the Advanced UI.
2. Click the **Setup** tab, and choose **Standby** from the list of setup pages.
3. Look at the status next to each configured partner. Possible states include:

State	Description
init	Initial state after configuring the partner IP address and enabling Standby
conn	Beginning the establishment of a connection with the partner
negotiate	Negotiating the connection
sync	Synchronizing partner's configuration and other data
running	A connection has been successfully set up with the partner.
inactive	Standby is disabled on current PacketShaper.

As PacketShaper sets up a connection with a partner, the state will quickly cycle through the stages of *init*, *conn*, *negotiate*, *sync*, and *running*. *Running* is the state you will see when the connection has been successfully established and the partner (and its link) is up. If the partner PacketShaper, or the switch or router it is connected to, goes down, the state alternates between *init* and *conn* as the active PacketShaper polls the partner and tries to reestablish a connection.

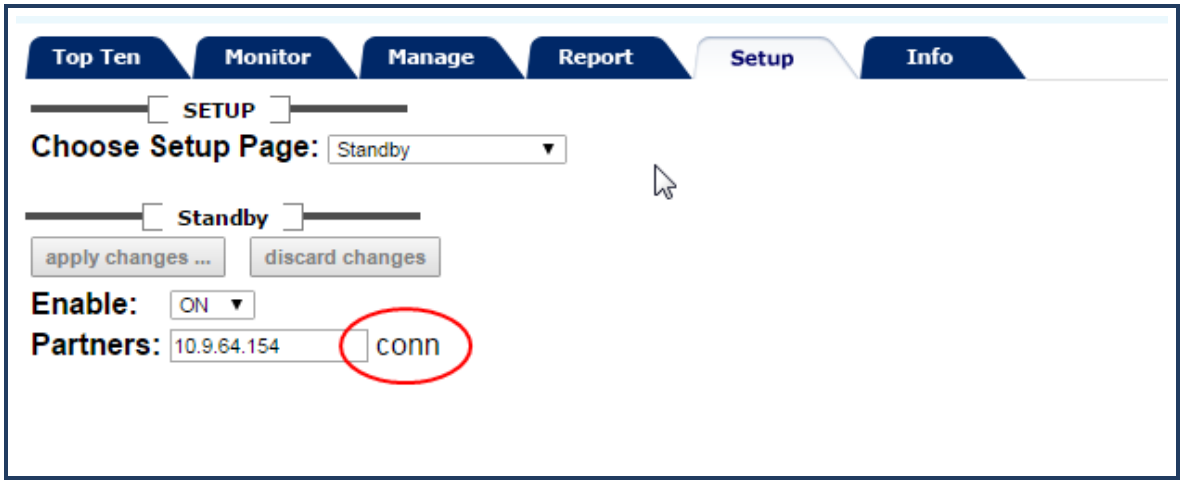

4. After a connection is set up with a partner, locate the additional tab that appears in the Advanced UI. This tab is labeled with the IP address of the partner PacketShaper.

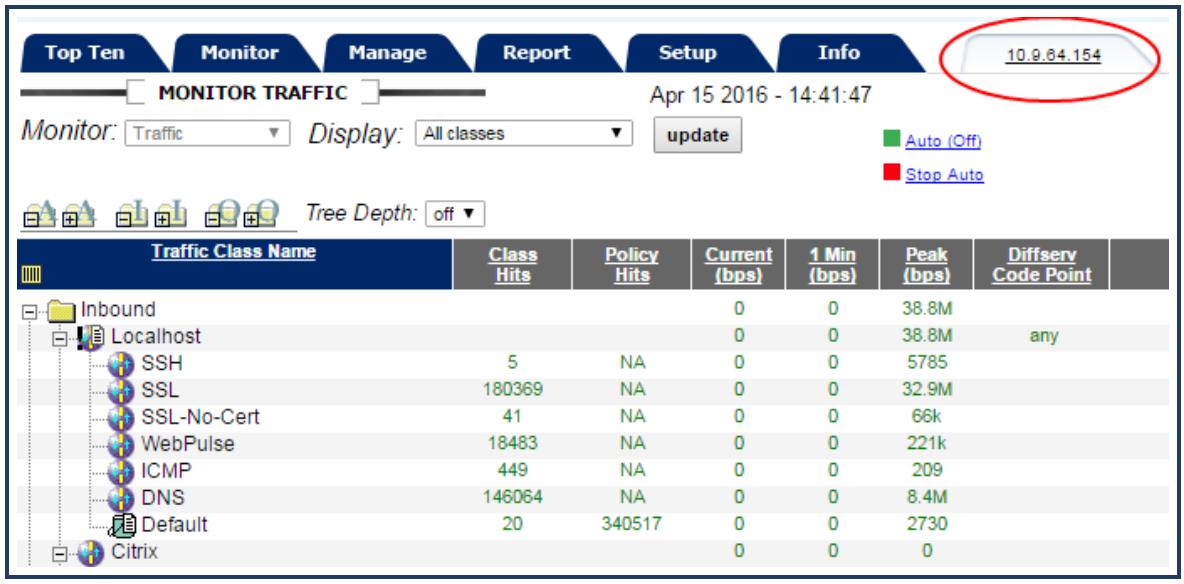

5. Click the partner tab to view the Monitor Traffic page of the partner PacketShaper. Make note of the following:
   - This page has a similar layout to the Monitor tab but it doesn't have active links to classes, partitions, and policies.
   - Use this tab to monitor traffic of the standby partner. As traffic goes through the partner unit, the statistics automatically update. Or, to refresh the screen, click **update**.
   - This page automatically reflects any traffic tree configuration changes made on the unit (such as new classes, modified partitions and policies). But you have to log in to the partner unit to make these changes.
6. To view reports of a standby partner:
   a. Click the **Report** tab. The default report is Network Performance Summary of the "self" PacketShaper.
   b. From the **Standby Partner** list, choose the IP address of the partner. The partner's Network Performance Summary report is displayed.

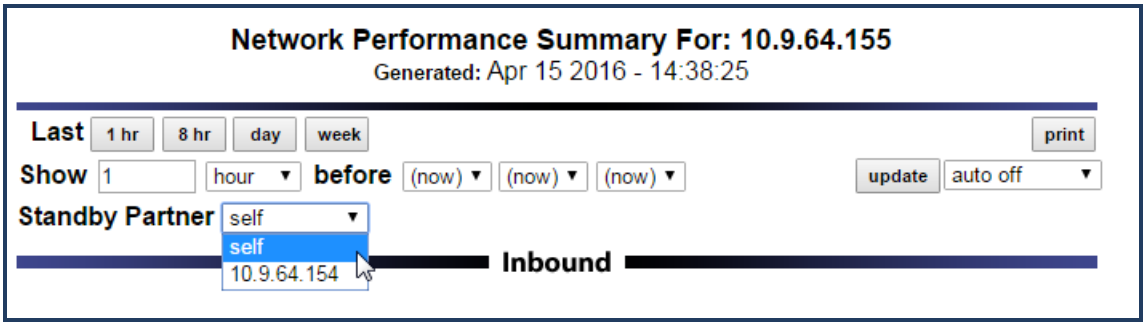

   c. You can also view partner reports of the **VoIP Summary** and **IPv6 Summary**.
7. Repeat the above steps for each partner in the cluster.

Standby Partner Failure

When a switch or router running redundancy protocol detects a link failure, it re-routes traffic to its alternate path. The PacketShaper in this alternate route will then take over traffic management responsibilities for the failed node, monitoring and shaping traffic that previously was managed by the PacketShaper in the disabled network path. This transition is seamless and transparent to users and administrators. The active node continues to manage its own traffic, in addition to managing the traffic on the disabled link.

Here's what happens when a PacketShaper takes over management of a failed partner:

- The partner monitor tab (labeled with its IP address) for the failed unit disappears from the Advanced UI of the active PacketShaper(s) it has partnered with.
- The Standby setup page in the active PacketShaper no longer shows *running* for the failed partner. Instead, the status alternates between *init* and *conn*, as the PacketShaper continues to poll the failed partner every five seconds.

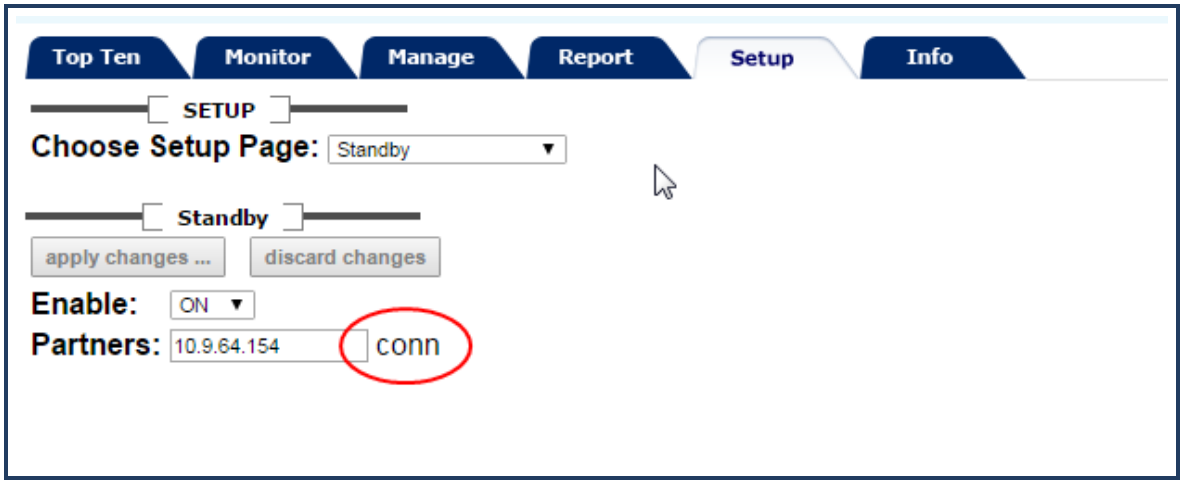

- The Monitor tab in the active PacketShaper contains the traffic tree of the failed partner. The failed partner's inbound tree is at the bottom of the Inbound branch of the active PacketShaper, and the outbound tree is under the Outbound branch. The failed partner's tree is in a folder labeled with *SB_* and the appliance serial number (for example *SB_031400000*).

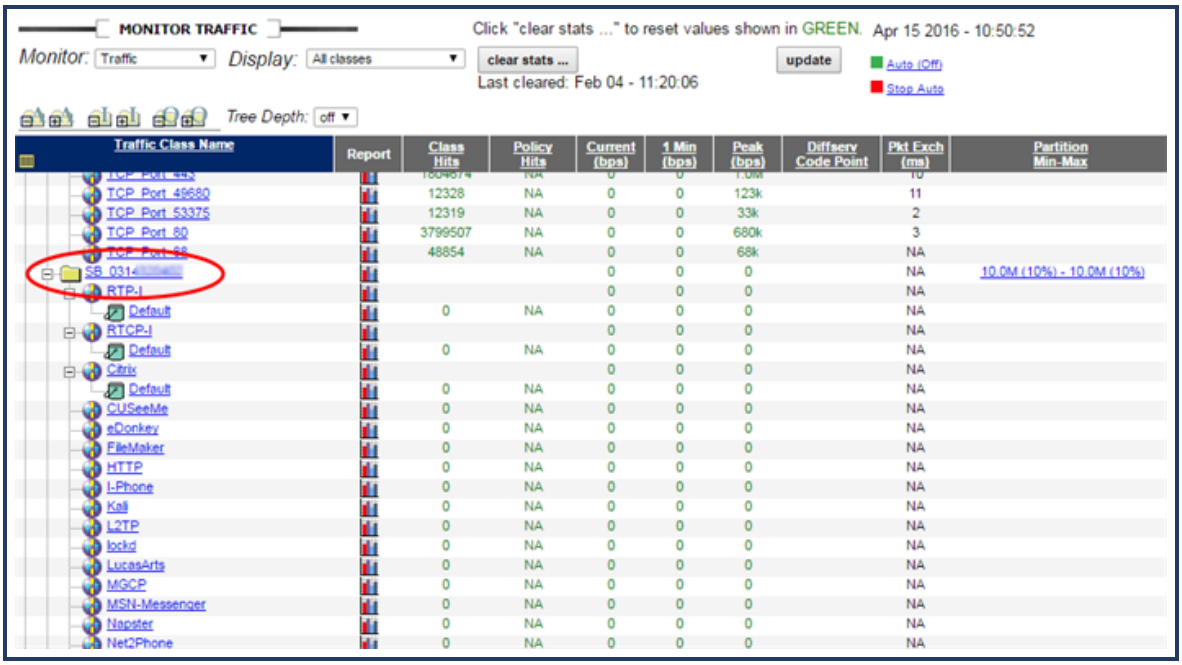

- While an active PacketShaper is taking over the traffic management responsibilities of a failed partner, the active unit uses the combined traffic tree to classify all flows that go through the appliance (which includes the traffic that has been re-routed to the active PacketShaper from the failed partner). It uses using its traditional methods of classification. For each passing traffic flow, PacketShaper traverses the traffic tree, starting at the top, to look for a matching traffic class. It compares each class' matching rules with the characteristics of the traffic flow. If they don't match, the search continues, going to the next class on the list (the next sibling) until it finds a matching traffic class.

  Because the failed partner's tree is merged below the active partner class tree, traffic previously controlled by the failed partner may very well get classified into a class in the active PacketShaper's traffic tree, and controlled by its policy or partition; only if the flow doesn't match any class in the active PacketShaper traffic tree will it get classified into a class in the failed partner's tree (under SB_SerialNumber folder). If you design unique class trees for each PacketShaper based on IP/subnet and if the classes are not overlapping between different PacketShapers, all re-routed traffic from the failed PacketShaper would hit the SB_SerialNumber class tree that has been merged from the failed PacketShaper.

- The **`setup standby show`** CLI output has a *Down* column that indicates the number of times the partner has gone down.

```
PacketShaper# setup standby show

Standby Status: ON

Standby Configuration:
ID  Serial No       IP Address      State Down   Uptime   Sent Bytes   Recv Bytes  Conns
 0 0314[illegible]      10.9.64.154       conn    3        0            0    203448056      0*

PacketShaper#
```

- When standby partners are down, the active and standby units share available bandwidth. Ten percent of available bandwidth is allocated to each standby device that is down. For example, if one standby partner is down, the active system uses 90 percent of the available bandwidth and the failed device gets 10 percent.

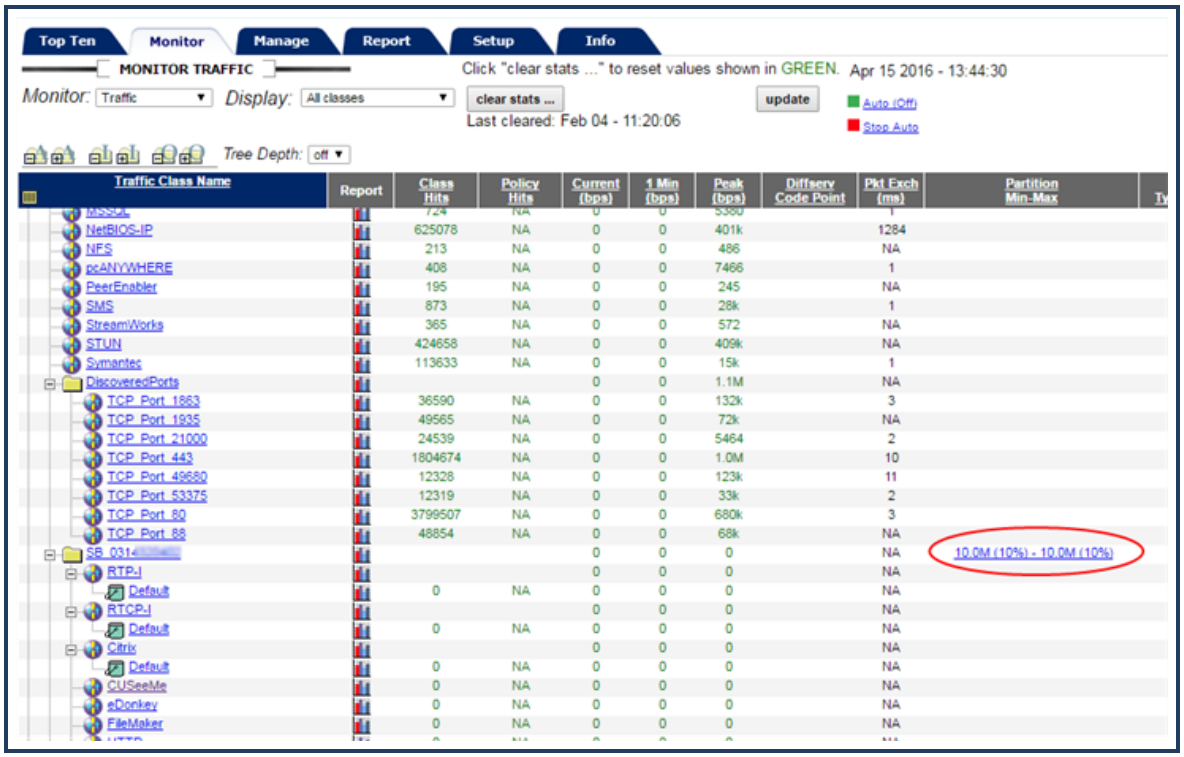

While the partner is down, you can modify its configuration from the active PacketShaper. This allows you to make policy and partition changes that might be necessary while its being managed by the active PacketShaper. For instance, to add a policy to a traffic class on the failed standby partner:

1. In the Advanced UI of the active PacketShaper, click the **Monitor** tab.
2. Under Inbound or Outbound, locate the folder of the standby partner (for example SB_031400000).
3. Locate the traffic class of interest and click the class name. The class attributes display in the Manage tab.
4. Select **policy > add** to create a policy.

> **Caution:** Any configuration changes you make to the failed standby partner will not be permanently saved to the configuration; once the partner is back up, it will return to its original configuration.

Returning to Active Status

The active PacketShaper polls the standby partner every five seconds to check its status. When the standby partner is running again:

- The partner tab will redisplay in the Advanced UI.
- The Standby page will show a *running* status next to the partner IP address.

As soon as the other link is back up, traffic will no longer be re-routed to the alternate route. Each PacketShaper will resume traffic management of its own link.

Disable Standby

To disable the standby feature, you need to turn standby off on each partner in the cluster.

1. In the Advanced UI, click the **Setup** tab.
2. From the **Choose Setup Page** list, select **Standby**. The Standby setup screen is displayed.

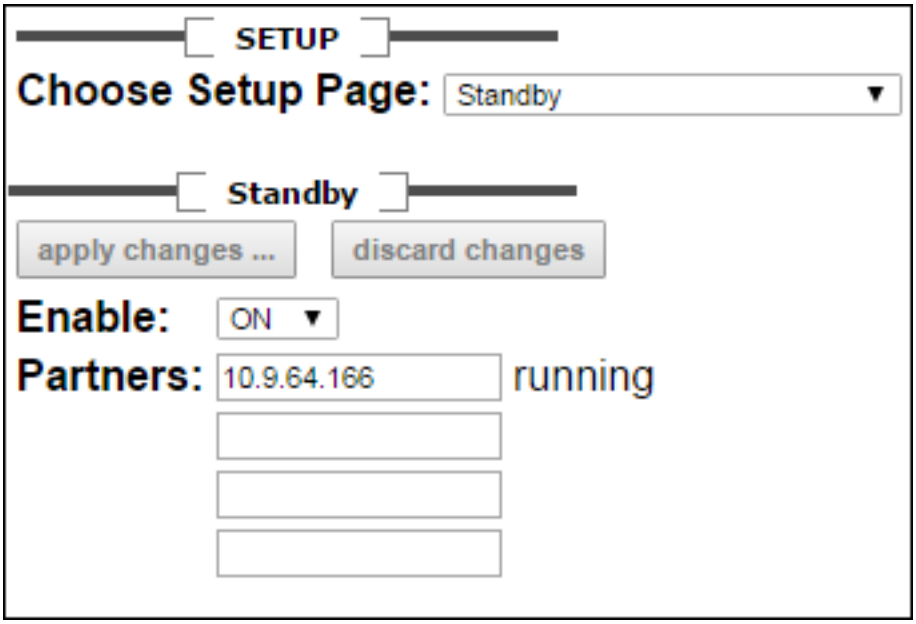

3. For **Enable**, choose **OFF**.
4. Click **apply changes**.
5. Repeat the above steps for each partner in the high availability cluster.
6. Re-enable bypass on the PacketShapers. Enter the following CLI command to set bypass to closed: **setup bypass close all** .

Caution: If you fail to re-enable bypass, the PacketShaper will not bypass traffic when it is powered down or in a failed state. In a non-standby configuration, all WAN traffic would be blocked.

Configure a High Availability Topology for Access-Link Monitoring

The high availability feature is also known as *access-link monitoring*. This feature is described fully in the " Access-Link Monitoring Overview " on page 51.

To configure high availability, you need to:

1. Define the access routers.
2. Define the WAN link interfaces.
3. Enable high availability by selecting a mode: basic or advanced.

Define High Availability Access Routers

To define a high availability access router:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **High Availability**.

3. In the *High Availability* setup page, click **add router**. The *Add Router* page is displayed.

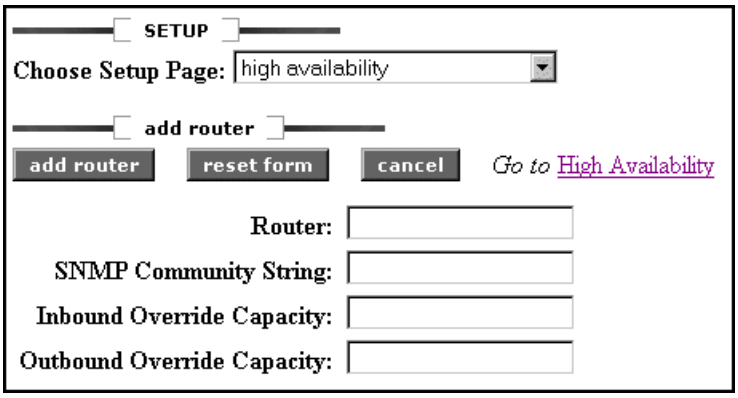

4. Define the router:

Field	Description
Router	IP address or DNS name of the access router
SNMP Community String	SNMP community string (password) for the router

5. (Optional) In the **Inbound Override Capacity** and **Outbound Override Capacity** fields, define the maximum throughput that is expected to pass through the router. Data rates are specified as integer bits per second, followed by a K (thousands), M (millions), or "G" (billions).

> **Note:** If you leave the override fields blank, PacketShaper uses the sum of the interface speeds to determine the WAN access line capacity for a router. However, in a configuration with multiple WAN interfaces on a router, you may want the inbound and outbound rates for the router to be less than the sum of the interface speeds, if you don't expect to get perfect load balancing between the interfaces. In this type of situation, you might want to configure override values. If both interfaces are up, PacketShaper would use the override values for the router when calculating the WAN access line capacity available for the router. If one of the interfaces goes down, PacketShaper would use the capacities configured for the active interfaces.

6. Click **add router**.

7. Add the WAN link interfaces for this router.

8. Repeat the above steps for additional access routers.

9. To see the routers and interfaces listed in the *High Availability Configuration* table, click **Go to High Availability**.

Define WAN Link Interfaces

To identify the WAN link interfaces on each high availability access router:

1. Make sure you have defined the router. The router's name and settings should be listed in the High Availability Configuration table on the *High Availability* setup page.

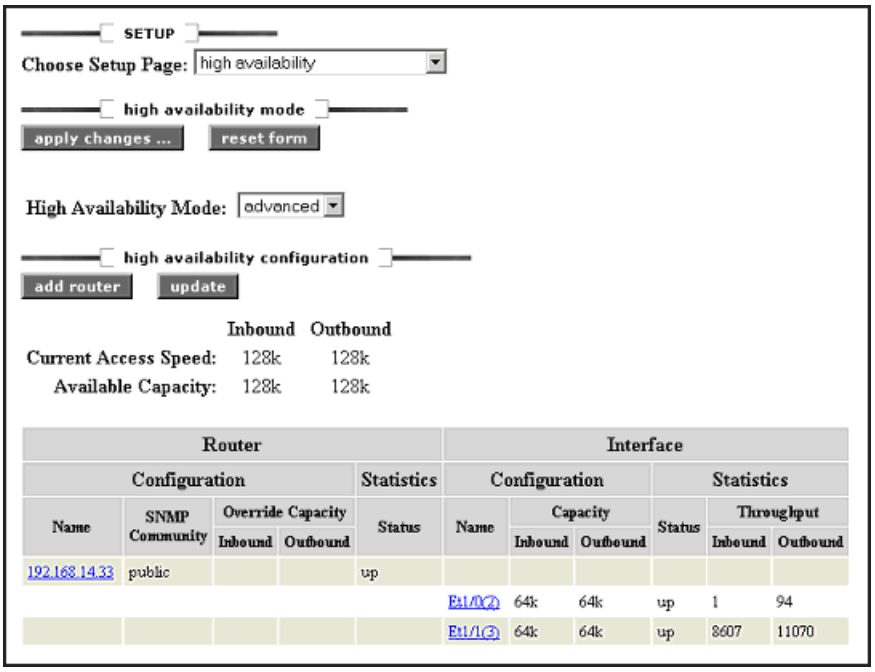

2. In the *High Availability Configuration* table, click the router's name.

3. Click **add interface**. The *Add Interface* screen is displayed.

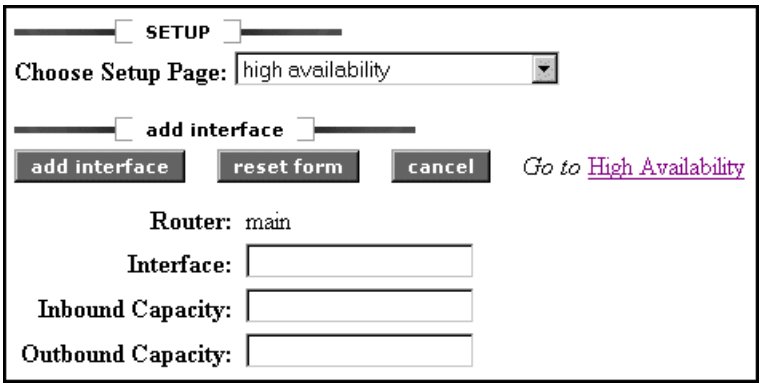

4. Define the interface:

Field	Description
Interface	Name (ifName) or index number (ifIndex) that identifies the interface. Examples of interface names are *ethernet 3/1* and *serial 0/1*. It is recommended that you identify the interface by name, not index, because ifnames are unique and persistent while index numbers can change dynamically. If you are using Cisco IOS v12.1 or above and have configured the router to make the ifindex persistent, you can safely identify the interface by index number. Note that ifname was not available in Cisco IOS before v11.1. Instructions for finding the ifName and ifIndex values for Cisco router interfaces
Inbound Capacity	Maximum inbound and outbound throughput that is expected to pass through the interface. Data rates are specified as integer bits per second, followed by a K (thousands), M (millions), or G (billions).
Outbound Capacity	

5. Click **add interface**. The interface settings are listed in the Interfaces table.

6. Repeat steps 3-5 above for each interface on the router.

7. When you're finished adding interfaces for the router, click **Go to High Availability**. The settings for each interface are listed in the *High Availability Configuration* table.

8. To define interfaces for other routers, repeat all the steps above.

Adding an interface will increase the router's available bandwidth unless you have set override values for the router. The lowest value (override versus sum of interfaces) takes precedence. For example, suppose a router has two 400K interfaces and you have set an override of 600K. The override takes precedence because it is the less than the sum of the interfaces (800K). Be aware that if you add another 200K interface, the override will continue to take precedence (in other words, the router's available bandwidth will still be 600K), even though it may not be appropriate. Make sure that you adjust your override after adding a new interface.

Enable Access-Link Monitoring

> **Note:** Leave the **High Availability Mode** field set to **off** until you have defined all your routers and interfaces.

After you have defined your access routers and their interfaces and are ready to enable access-link monitoring, follow these steps:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **High Availability**. The *High Availability* screen is displayed.

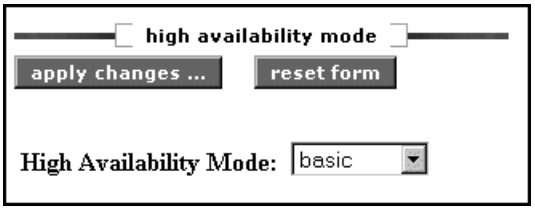

3. In the **High Availability Mode** field, select one of the following:

- **basic** When the basic mode is enabled, PacketShaper polls the configured router(s) every 30 seconds to assess the status (link up or link down) of the WAN link interfaces. If a link goes down, PacketShaper will automatically adjust the total available capacity by subtracting out the capacity of the down link. As part of this process, it will adjust the access link size and resize Inbound and/or Outbound partitions to reflect the available bandwidth.

- **advanced** Advanced mode includes all the functionality of basic mode, plus offers link overload protection. PacketShaper will use SNMP polling to access the actual throughput of each configured WAN link interface. If an interface approaches its configured capacity, PacketShaper will pace the traffic sent to the router to prevent overloading any interface. This pacing will also greatly reduce the number of retransmissions.

3. Click **apply changes**.

> **Note:** In order for the partition size adjustments to take effect, traffic shaping must be enabled.

After high availability is enabled, you will see additional information in the *High Availability Configuration*.

Item	Description
Current Access Speed	The Current Access Speed is the total of all interface capacities that are active (up). A router's capacity is determined by the values set for the override capacity or by summing all the interfaces' capacities (if no override has been set). When high availability is enabled and a link becomes inactive, the Current Access Speed will reflect this reduction of available bandwidth (that is, the inactive link's capacity will be subtracted out, assuming it is less than the override value).
Available Capacity	The Available Capacity is the total bps available based on the values configured for the interfaces and routers. It is the sum of the routers' capacities (without regard to whether its interfaces are up or down).

Item	Description
Statistics	**Router Status**: Indicates whether the router is up or down **Interface Status**: Indicates whether the interface is up or down **Interface Throughput**: The actual throughput (based on SNMP polling) measured in bits per second (bps); this information is shown only in advanced mode

View the High Availability Configuration

The high availability configuration is displayed in a table on the High Availability setup page.

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **High Availability**.

Alternatively, you can click the **Go to High Availability** link that is available on the various high availability setup pages (such as *Edit Router* and *Edit Interface*).

The table below describes the items in the *High Availability Configuration* table.

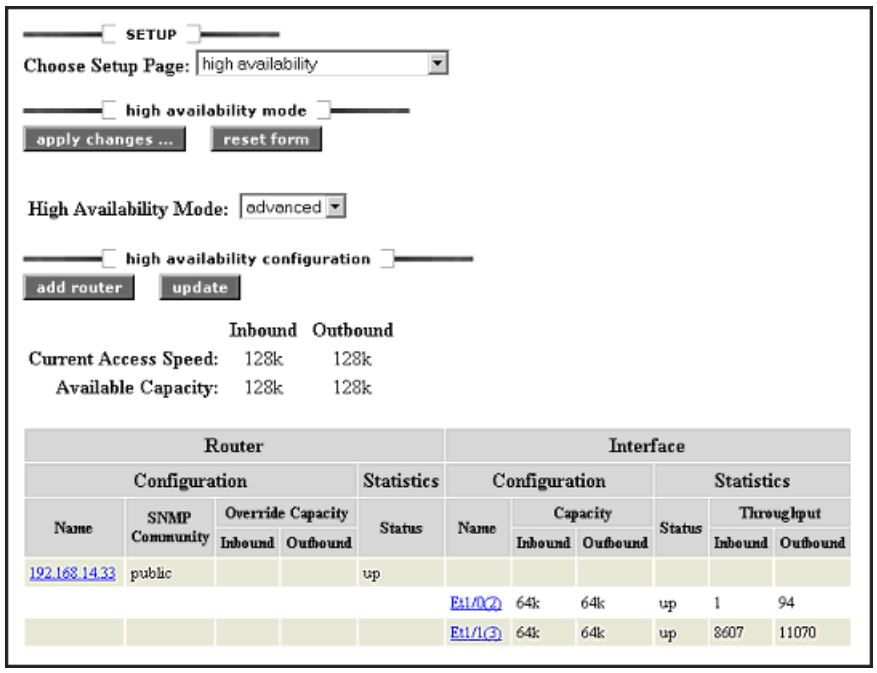

Item	Description
High Availability Mode	The current high availability mode (basic, advanced, or off). See " Enable Access-Link Monitoring" on page 736 for more information. **Note:** Different information appears in the table based on the mode selected. If you haven't enabled high availability or you haven't defined your routers and interfaces, you will not see all the information described below.
Current Access Speed	The total of all interface capacities that are active (up). A router's capacity is determined by the values set for the override capacity or by summing all the interfaces' capacities (if no override has been set). When high availability is enabled and a link becomes inactive, the Current Access Speed will reflect this reduction of available bandwidth (that is, the inactive link's capacity will be subtracted out, assuming it is less than the override value).
Available Capacity	The total bps available based on the values configured for the interfaces and routers. It is the sum of the routers' capacities (without regard to whether its interfaces are up or down).

Item	Description
Router	The router's configured name (IP address or DNS name), SNMP community string, the override capacity (if set), and the router status (up vs. down). See Define High Availability Access Routers for more information. **Note:** Click the router's name to modify the router settings.
Interface	The interface's configured name and number (ifname and ifindex), the inbound and outbound capacities, and the interface status (up vs. down). If the interface name (ifname) is not found on the router, the Name column displays a zero after the name, for example *test(0)*. If the interface number (ifindex) is not found, the Name column displays Unknown, for example *Unknown(1234)*. If advanced mode is enabled, the actual bps throughput (based on SNMP polling) is listed for each interface. See Define WAN Link Interfaces for more information. **Note:** Click the interface name to modify the interface settings.

Modify High Availability Access Routers

If you need to change a high availability router's community string or override capacities, you can edit the router's configuration.

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **High Availability**. The *High Availability* screen is displayed.

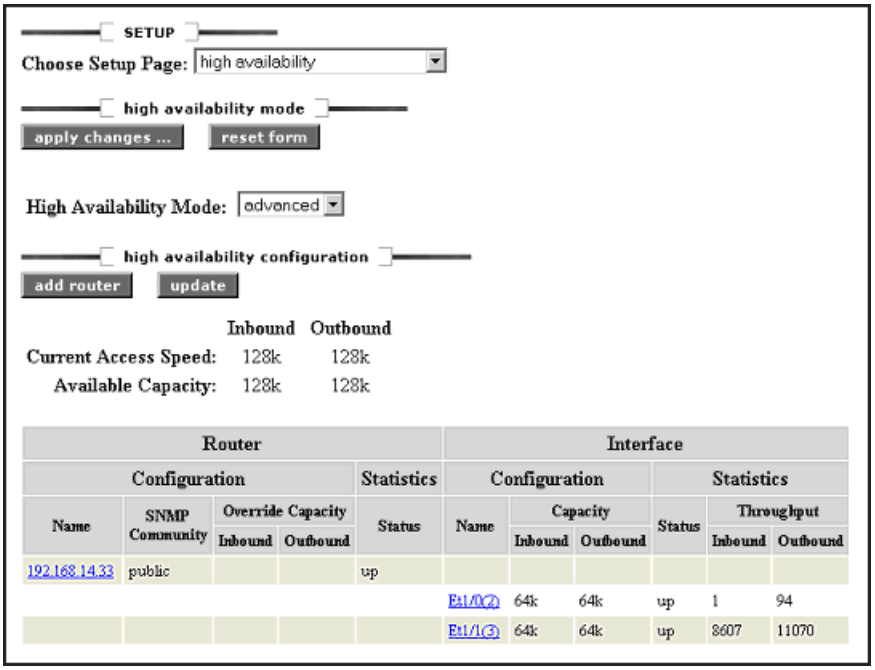

3. In the *High Availability Configuration* table, click the router's name. The *Edit Router* screen is displayed.

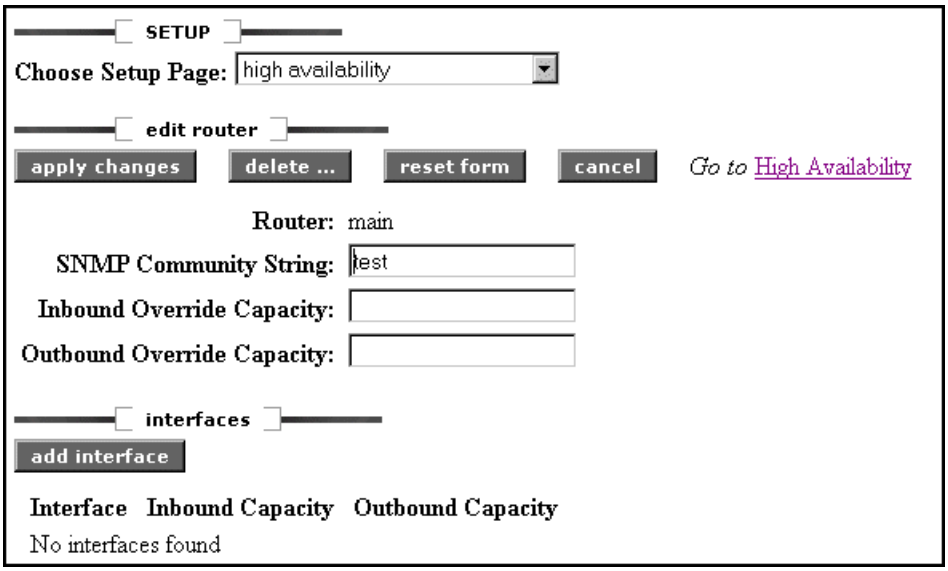

4. Modify the settings.

5. Click **apply changes**.

Modify High Availability Interfaces

If you need to adjust the access speeds of a previously-defined WAN link interface, there are two ways to modify the settings:

- In the High Availability Configuration table, click the interface name you want to modify.

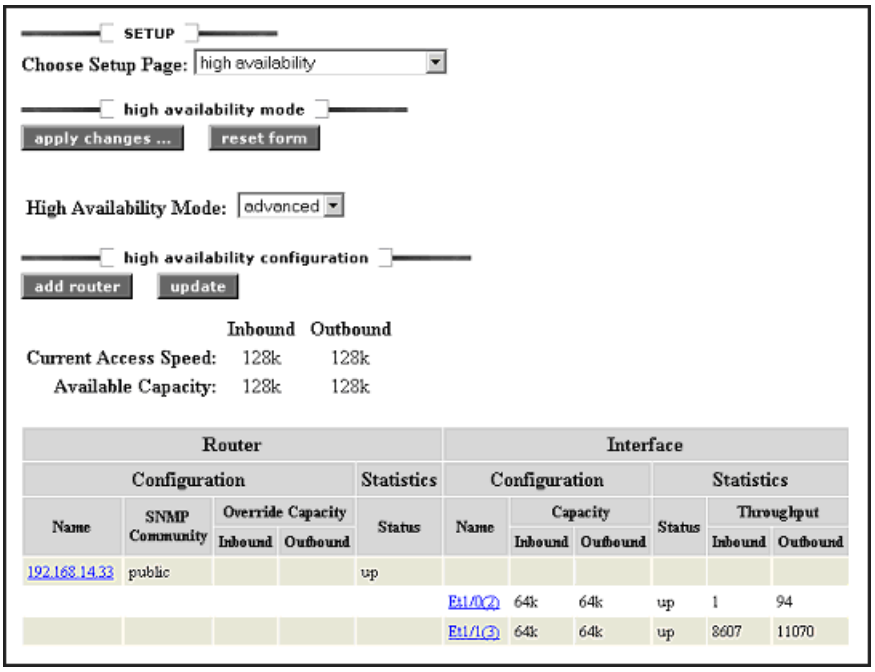

- When modifying the router settings on the *Edit Router* page, you can click **edit** next to the interface you want to modify.

After you make your changes, be sure to click **apply changes**.

Remove High Availability Access Routers

To remove an existing router from the high availability configuration:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **High Availability**. The *High Availability* screen is displayed.

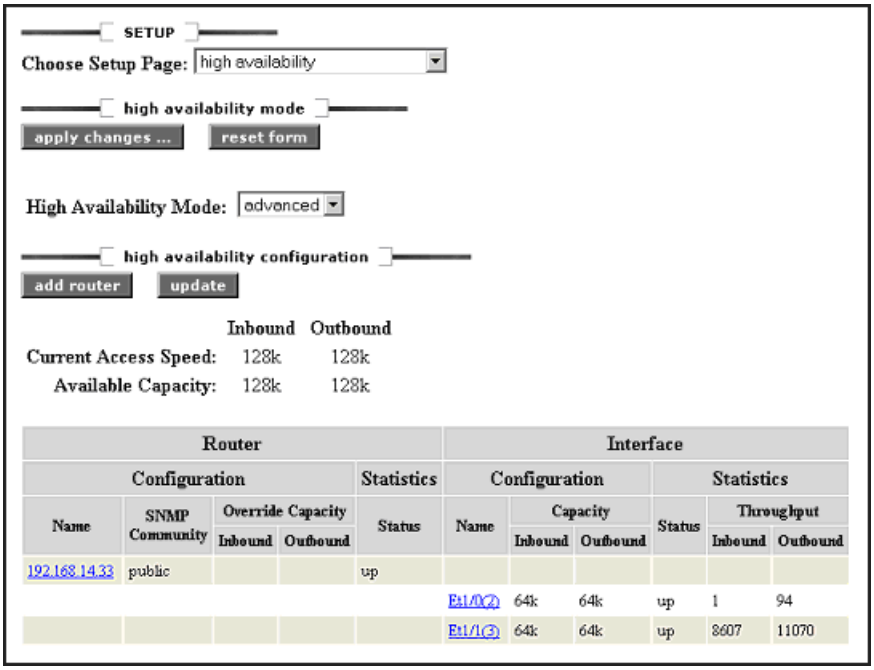

2. In the *High Availability Configuration* table, click the router's name. The *Edit Router* screen appears.

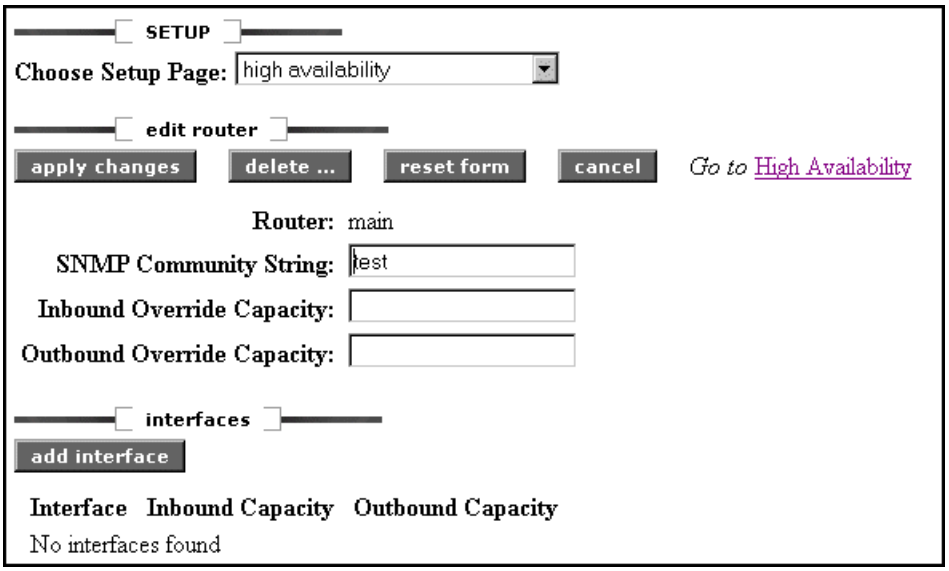

3. Click **delete**.

After you confirm, the router and its settings are removed from the *High Availability Configuration* table.

Remove High Availability Interfaces

There are two ways to remove an interface from the high availability configuration:

- In the High Availability Configuration table, click the interface name you want to remove and then click **delete**.

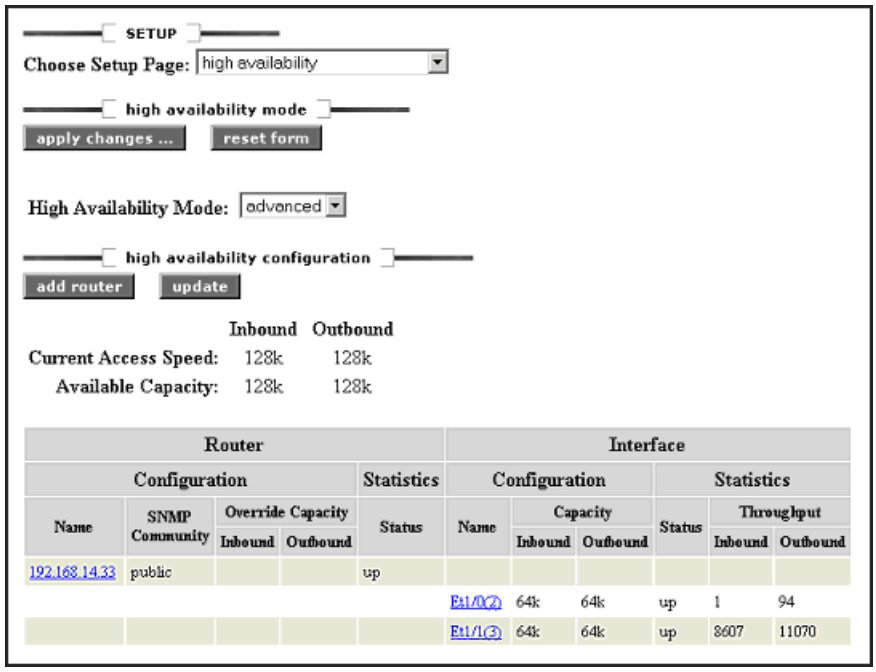

- When modifying the router settings on the *Edit Router* page , you can click **delete** next to the interface you want to remove.

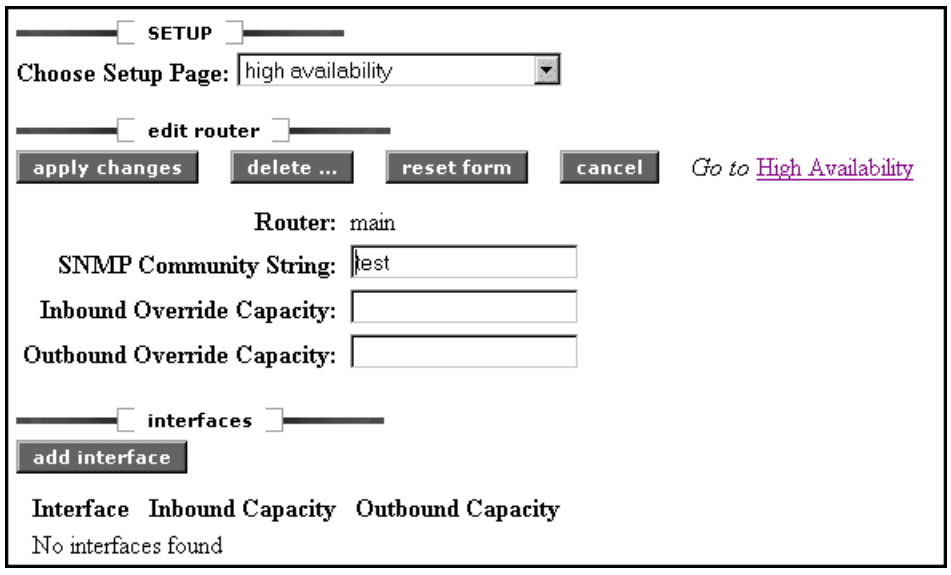

After you confirm, the interface and its settings are removed from the High Availability Configuration table. (Click **cancel** to see this table.)

Deleting an interface may reduce the router's available bandwidth, depending on the override value. For example, suppose a router has two 400K interfaces and you have set an override of 600K. If you then delete an interface, the router's available bandwidth would be reduced to 400K; the override would be ignored since it's greater than the sum of the router's interfaces.

Customer Portal Tasks

The customer portal feature offers service providers the ability to design custom network and application status pages that can be viewed by their customers. This section includes tasks for creating a login page, designing portal content, and creating customer portal accounts. For more information, see the " Customer Portal Overview" on page 76.

Create Customer Portal Accounts

To set up your customer portal accounts, you need to define each customer's login name, password, home directory, and an optional message that appears on the customer's portal page.

To add a customer:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, choose **Customer Portal Administration**. The *Customer Portal* screen appears.

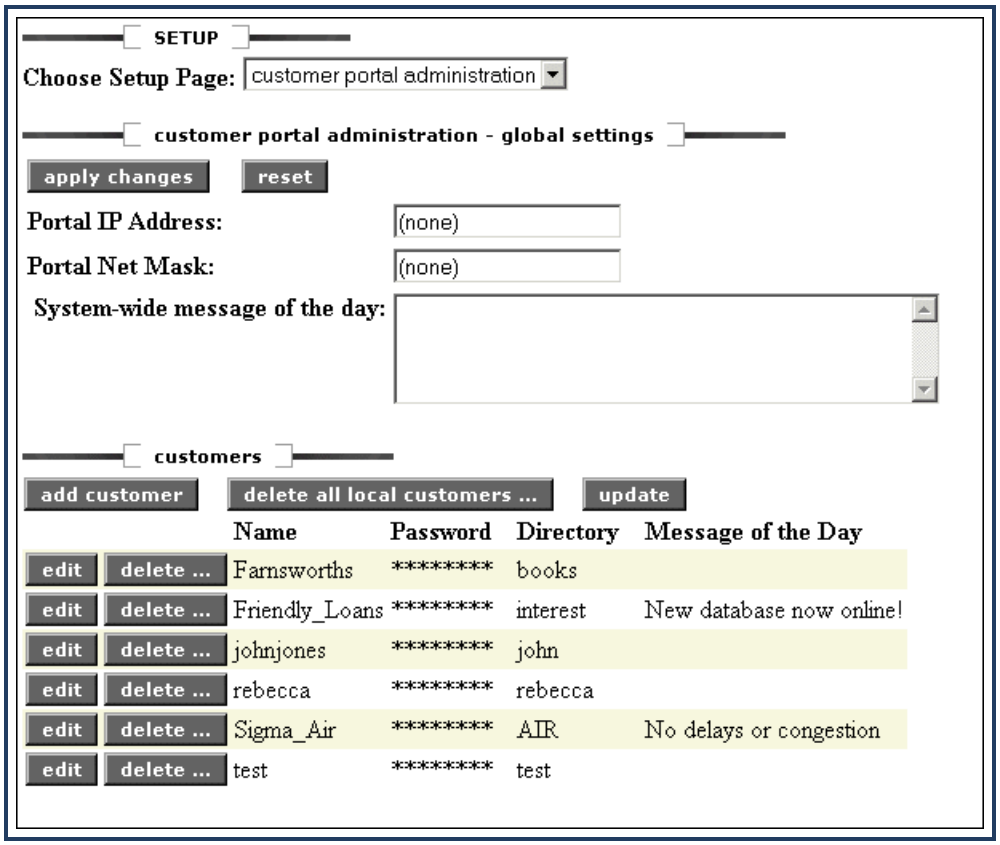

3. Click **add customer**. The *Add Customer* page appears.

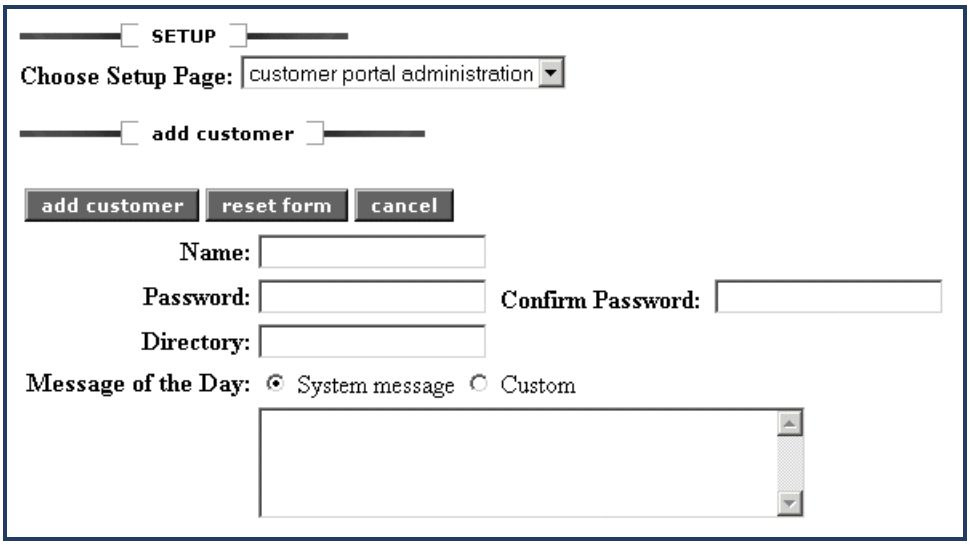

4. Enter the customer's information:

Field	Description
Name	The login name the customer will use; up to 32 characters long, use numbers, letters and underscores– spaces are not allowed. If you are using RADIUS or TACACS+ authentication, this name must match the user name entered in the RADIUS/TACACS+ server.
Password	The password for the customer account. If you are using RADIUS or TACACS+ authentication, you do not specify a password here– the customer portal will use the password specified for this user in the RADIUS/TACACS+ server.
Confirm Password	Type the same password a second time
Directory	The name for the customer home directory to be created in the path 9.258/customer/, maximum 8 alphanumeric characters. If you want multiple customers to access the same portal page, specify the same directory name for all of these users (for example, 9.258/customer/master).
Message of the Day	Select **System message** to use the system-wide message or select **Custom** and then enter a text message, up to 128 characters. This custom message can be displayed on the customer home page. (See also Set a System-Wide Message.)

5. Click **add customer**.

The Add Customer screen creates a home directory for each customer you configure, but these directories are initially empty. You must create an individual INDEX.HTM file and place it in the home directory of each customer. Or, if multiple customers will be accessing the same portal page, place the INDEX.HTM file in the master directory specified for these customers. See Design Customer Portal Pages.

Edit Customer Information

In the *Edit Customer* screen you can change the password, directory, or message for an existing customer. Note that you cannot change the customer login name.

To edit customer information:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, choose **Customer Portal Administration**. The *Customer Portal* screen appears.

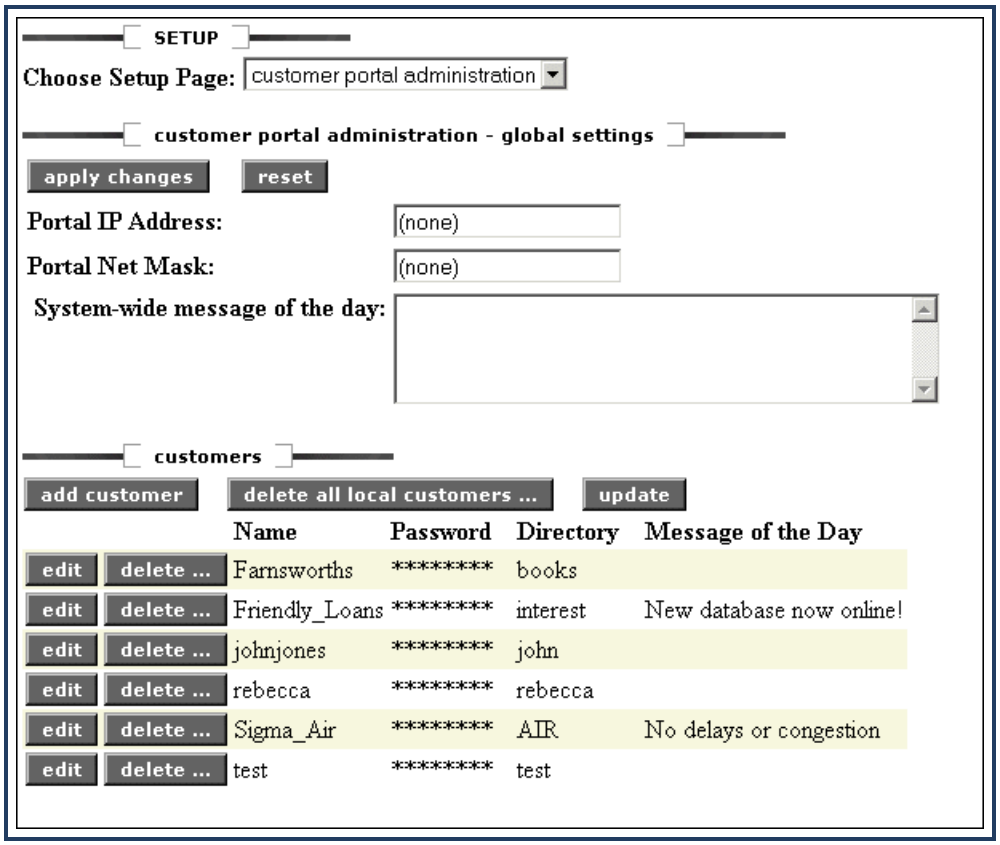

3. Click **edit** next to the customer whose information you want to modify. The *Edit Customer* screen appears.

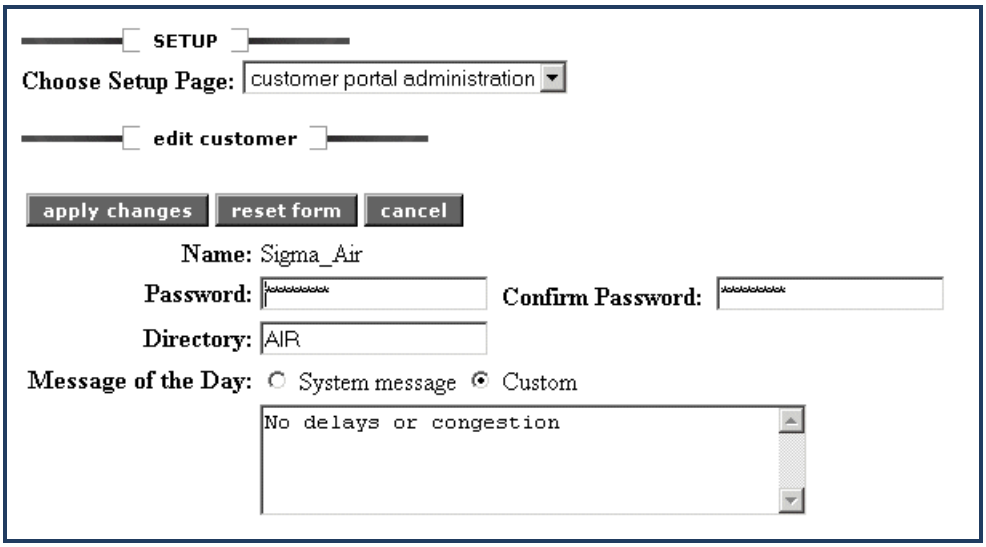

4. Edit the customer's existing information:

Field	Function
Password	The password for the customer account. If you are using RADIUS or TACACS+ authentication, you do not specify a password here– the customer portal will use the password specified for this user in the RADIUS/TACACS+ server.
Directory	The name for the customer home directory to be created in the path 9.258/customer/, directory names are limited to 8 alphanumeric characters
Message of the Day	Select **System message** to use the system-wide message or select **Custom** and then enter a text message, up to 128 characters. This custom message will be displayed on the customer home page. (See also Set a System-Wide Message.)

5. Click **apply changes**.

Remove Customer Accounts

You can delete customer accounts one at a time or all the accounts at once.

To delete customer accounts:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, choose **Customer Portal Administration**. The *Customer Portal* screen appears.

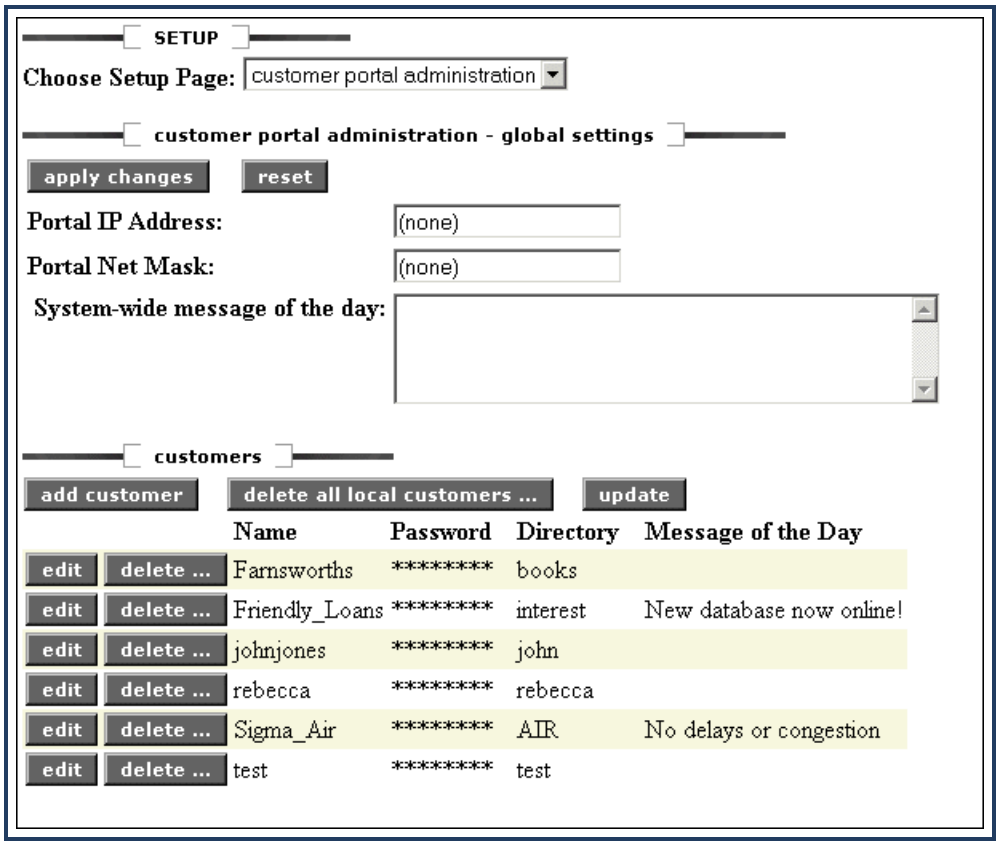

3. To delete one customer account, click **delete** next to the customer whose account you want to delete.

or

To delete all customer accounts, click **delete all local customers**.

4. Click **OK** to confirm the deletion.

Set a System-Wide Message

The *message of the day* is a text string that appears on the customer portal pages. The system-wide message is used for any customer accounts that do not have a custom message configured.

To set a system-wide message:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, choose **customer portal administration**. The *Customer Portal* screen appears.

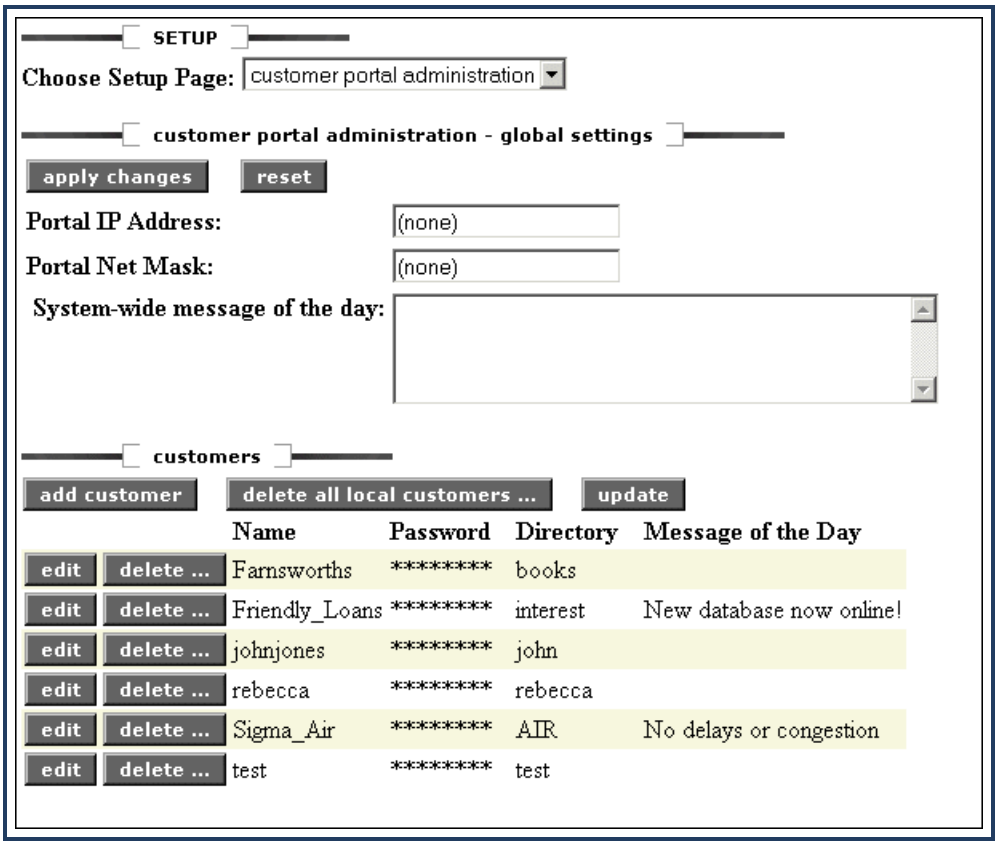

3. In the *Customer Portal Administration* screen, click in the **System-wide message of the day** field.

4. Type the text message, up to 128 characters. The following characters are not allowed: ";\ & '

> **Caution:** If you use the carriage return or **Enter** buttons on your keyboard to create multiple lines of text, make sure that each line has a printing character.

5. Click **apply changes**.

Download Sample Portal Files

Symantec offers sample login and customer pages you can download and customize for your own use.

1. Log in to BTO (https://bto.bluecoat.com) using your BlueTouch Online credentials.

2. Go to the **Downloads** tab and navigate to the PacketShaper page for your model.

3. Choose the PacketShaper software version you are using.

4. Download the portal zip file (for example, *PacketShaper_11.7.1_portals.zip*).

5. Extract the files to your computer for editing.

6. When you have finished editing the files, use the File Browser to upload them to the locations specified below.

The customer portal zip file contains the following files:

File	Description
The following files are different types of sample customer portal pages:	
ex1.htm	A sample HTML customer page. See Design Customer Portal Pages for details on using this sample file.
glast24.htm	A sample HTML page with graphs of the preceding 24 hours network performance for: ■ Inbound throughput ■ Outbound throughput ■ Inbound efficiency ■ Outbound efficiency To use this file, either place a customized version in each customer's directory or put an edited version in the 9.258/customer directory for all customers to access.
glast60.htm	A sample HTML page with graphs of the preceding 60 minutes network performance for: ■ Inbound throughput ■ Outbound throughput ■ Inbound efficiency ■ Outbound efficiency To use this file, either place a customized version in each customer's directory or put an edited version in the 9.258/customer directory for all customers to access.
owner1.htm	A sample customer portal page that illustrates how to use the SEARCH facility in order to have one portal page serve many users. This page searches the traffic class' Owner field for the name of the logged-in portal user. To use this file, place an edited version in the 9.258/customer/*<shared>* directory, where *<shared>* is the name of the "community" customer directory. (**Note:** The *<shared>* directory name must be specified as the directory name for each customer that will be using this portal page. See Create Customer Portal Accounts.)
notes1.htm	Similar to owner1.htm, except this page searches the Comment field.
toptalk.htm	A sample HTML file that shows how to display Top Talkers and Top Listeners data in a portal page.

File	Description
The following files are used for the login page:	
blank.gif	Image file used by the sample login page. Place this file in the 9.258/customer directory.
lginbtn.gif	Image file of the login button on the sample login page. Place this file in the 9.258/customer directory.
libmd5.js	JavaScript that supplies an MD5 method (an algorithm used to create digital signatures) for the authentication procedure. Place this file in the 9.258/customer directory.
login.htm	Sample HTML login page. See Customize the Login Screen.

Note: If you view the sample HTML files directly in a web browser, instead of fetching them from the PacketShaper software, error messages and text bracketed by percent signs may appear. This happens because the web browser does not know how to process the special tag engine instructions in the Symantec sample pages.

Customize the Login Screen

All customers accessing the customer portal see the same initial login screen. When a customer accesses the URL http://*<PacketShaper-ip-address>*/customer/ or http://*<portal-ip-address>* a login screen appears, prompting the customer to enter a name and password.

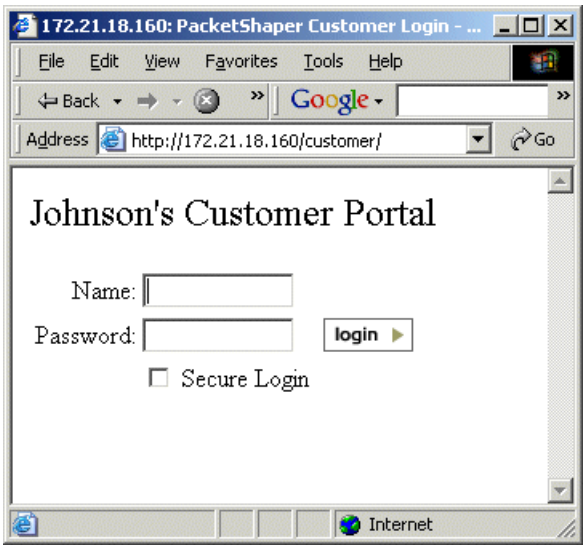

The customer then enters the preassigned login name and password. After the customer clicks the **login** button, the index.htm file is presented from the appropriate customer directory beneath 9.258/customer/.

To design the login screen:

1. Edit the sample login page you downloaded (LOGIN.HTM).

2. Customize this page, for example by adding graphics or headers to identify the service provider.

3. Use the File Browser to upload your edited LOGIN.HTM file to the 9.258/customer directory on the PacketShaper.

Design Customer Portal Pages

To design a customer portal page:

1. Edit the sample customer page you downloaded (EX1.HTM). See " Download Sample Portal Files" on page 754.

2. Determine the traffic classes for which users will need to see statistics and specify these classes in the portal page. See "Select Traffic Classes for the Portal Page" on page 1091.

3. Determine the measurement variables to include on the portal page. See "Measurement Variables Available to Portal Pages" on page 1103.

4. Customize the page to display the desired statistics of the desired classes to the users of the customer portal. See "Display Statistics on the Portal Page" on page 1097

5. Include custom information such as the customer name and message of the day on the portal page. See "Display Custom Information on the Portal Page" on page 1093.

6. *Optional:* Include graphs on the portal page. See "Include Graphs on a Portal Page" on page 1094.

7. *Optional:* Link to other HTML pages. See "Link to Other HTML Pages" on page 1096.

8. Save your customized file as INDEX.HTM.

9. Use the File Browser to upload your customized INDEX.HTM file to the customer's directory in the 9.258/customer directory on the PacketShaper for example, *9.258/customer/jones*. Or, if multiple customers will be accessing the same portal page, place the INDEX.HTM file in the master directory specified for these customers.

Examples

Example 1 (without graphs): Let's put it all together now into a sample page designed for Farnsworths, a retailer of books, music, and videos.

This sample has two basic objects:

- A table with aggregate throughput and efficiency data for /Inbound and /Outbound. For each entry, there is a square which is displayed as red, yellow, or green, depending on the value of the avg-bps variable.
- For the HTTP traffic class, a table of response times.

You can view a sample page with the source code presented as readable text. See "Sample Customer Portal Page" on page 1121.

Example 2 (with graphs): This example shows the throughput and network efficiency for both Inbound and Outbound over the preceding 60 minute period. See "Sample Customer Portal Page with 60 Minute Graphs" on page 1120.

Assign a Customer Portal Address

You can assign an IP address to the customer portal, allowing customers to directly display the customer portal login page using this address instead of the http://x.x.x.x/customer URL.

To assign a portal address:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, choose **customer portal administration**. The *Customer Portal* screen appears.

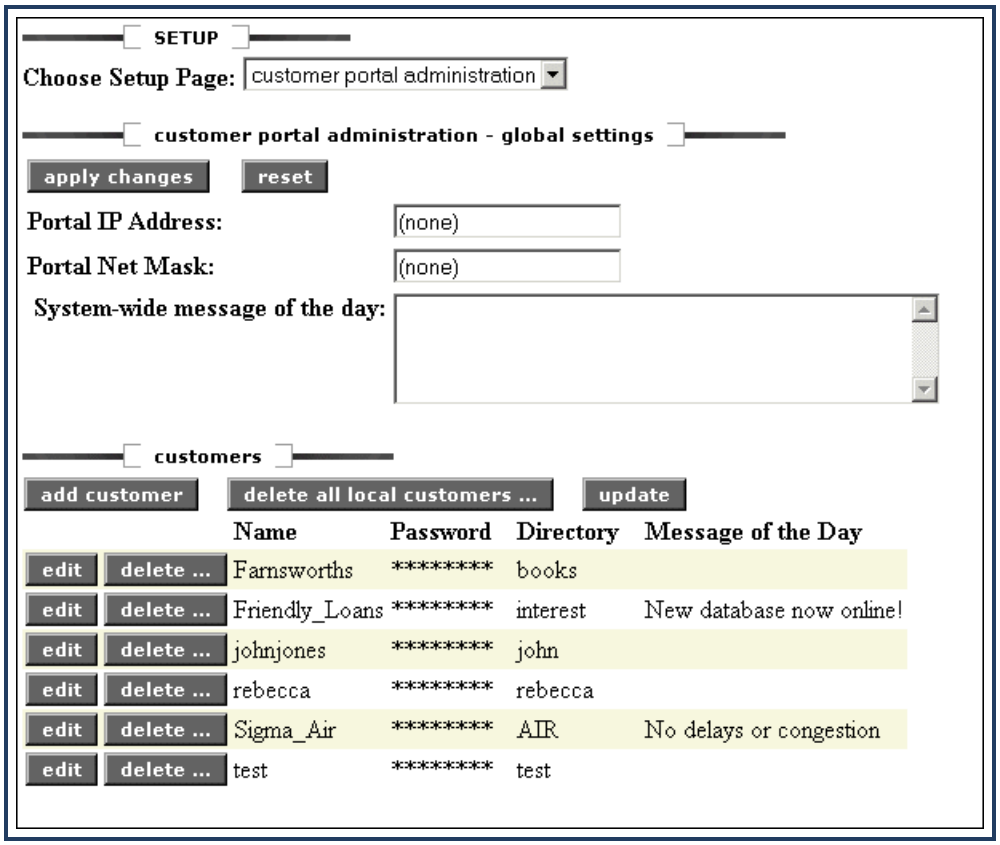

3. In the *Customer Portal Administration* section, click in the **Portal IP Address** field and type the IP address you want to assign to the customer portal. The address should not be the same as the management IP address.

4. Click in the **Portal Net Mask** field and type the subnet mask for the network where the unit resides.

5. Click **apply changes**.

Log Into the Customer Portal

After you have created your login screen, designed your portal page, and copied these files to the proper locations on the PacketShaper, you can log into the customer portal.

To log into the customer portal:

1. In a web browser, type **http://*<PacketShaper-ip-address>*/customer** where *<PacketShaper-ip-address>* is the management IP address of the PacketShaper. For example: *http://10.10.100.10/customer*

   or

   If you have assigned an IP address to the customer portal, type **http://*<portal-ip-address>*** where *<portal-ip-address>* is the address you assigned. For example: *http://10.10.100.11*

   You will then be presented with the login screen you created. For example:

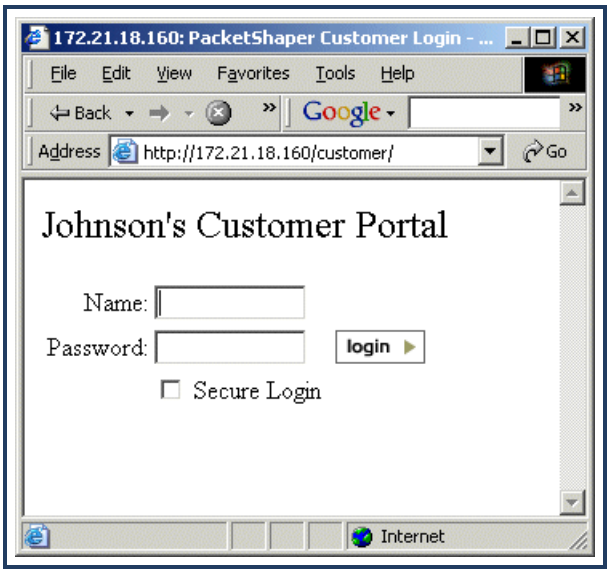

2. On the login page, enter the customer's login name and password that were specified when the customer accounts were created. (See Create Customer Portal Accounts.)

3. Click **login**.

   The customer portal page you designed is displayed.

Configure Captive Portal

PacketShaper's *captive portal* feature allows enterprises and service providers to present an Acceptable Use Policy (AUP) to guest users accessing the Internet in a web browser on a WiFi network. See "Captive Portal Overview" on page 70.

Create Custom AUP Pages

PacketShaper includes a sample AUP HTML page that you can use as is, or customize with company policy verbiage and logos. You can create a customized AUP page for each customer and/or site.

1. In the PacketShaper Advanced UI, go to the Info tab and open the File Browser.
2. Navigate to `9.258/customer/aup`.
3. Locate the sample HTML AUP page, aup_example.htm.

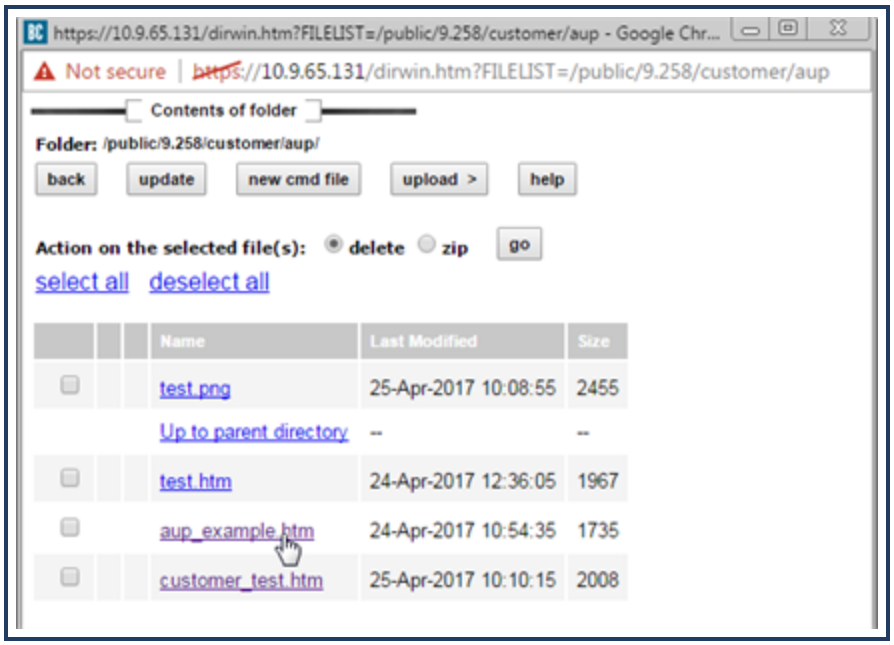

4. View the sample page and save it to your local workstation.

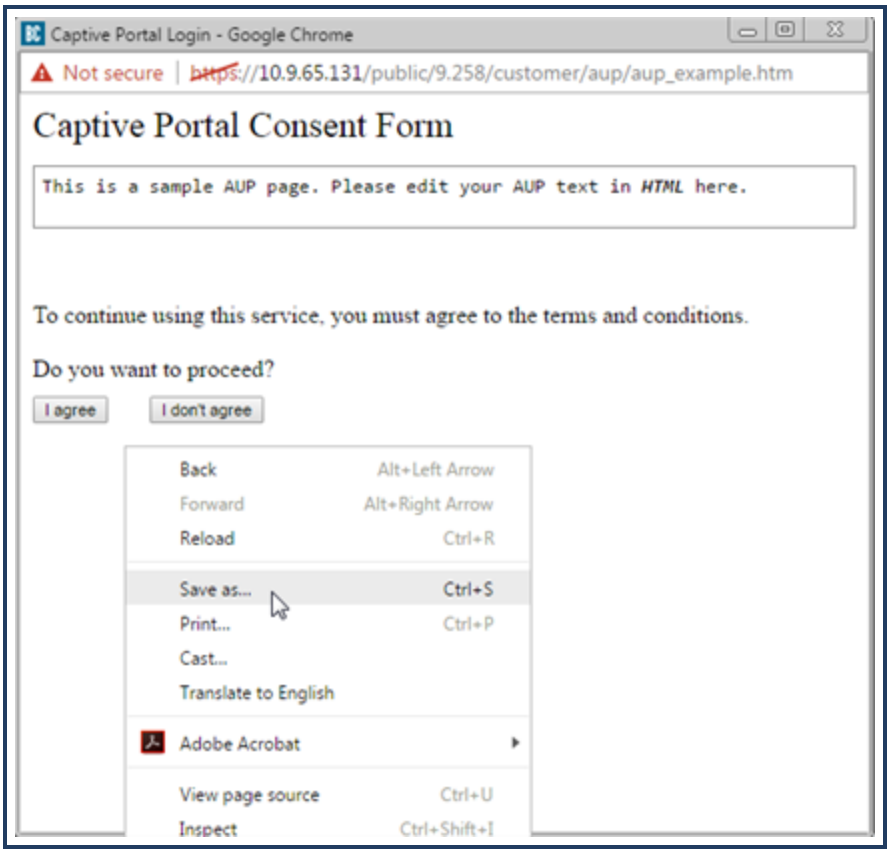

5. Open the sample page in a text or HTML editor.
6. Customize the page as desired, for each customer and/or site.

> **Caution:** Do *not* modify the section containing the **Agree** and **I don't agree** buttons.

7. To include a logo or other image, use the `<img src>` tag. For example:

```
<font size="+2">Test Corporation Guest WiFi</font>
<p><img src="test.png"></p>
```

8. Save each customized page with a name that identifies the customer or site ID. Make sure the file has a .htm extension (not .html).
9. Upload the HTML files and any referenced image files to PacketShaper's 9.258/customer/aup directory.
    a. Open the File Browser and navigate to `9.258/customer/aup`.
    b. Click **upload >**.
    c. Choose the file to upload.
    d. Click **upload**.
    e. Repeat for each file.

10. To verify the file, enter the following URL in the address line of your web browser: *<PacketShaper IP>*/customer/aup/*<filename>*. For example: *10.10.10.10/customer/aup/customer1_aup.htm*

Configure Captive Portal in the CLI

To set up the captive portal feature on PacketShaper, use the **tenant** commands in the command-line interface.

1. . Access the command-line interface.
2. Define the subnet to be managed on your network. This subnet needs to encompass all network segments in the guest network that you want to manage with captive portal. See "tenant set" on page 1384. For example:
   ```
   tenant set start-ip 10.100.0.0
   tenant set netmask 255.255.0.0
   ```
3. Define one or more captive portal IP addresses. This IP address is used internally to perform web redirection (for example, to redirect from the user's requested URL to the AUP page). The IP can be any valid address that is not in use by other hosts on the network. You can add up to eight captive portal IP addresses, with each one serving a different AUP (for example, customized to different customers). See "tenant set" on page 1384.

   ```
   tenant set captive add 1.2.3.4
   ```
4. Add one or more customers. Service providers can assign an ID to each customer. Even if your organization doesn't have multiple customers, you must create an ID so that you can then define the default AUP for the company. See "tenant add customer" on page 1377.

   ```
   tenant add customer id 1
   ```
5. Specify the name of the AUP file for each customer. See "tenant add customer aup" on page 1378.

   ```
   tenant add customer 1 aup "aup/customer1_aup.htm"
   ```
6. Define one or more sites for each customer. If a customer has multiple subnets (each with a different Acceptable Use Policy), you will create a unique site for each network segment. See "tenant add site" on page 1379.

   ```
   tenant add site id 1 2
   ```
7. Set the attributes for each site. At a minimum, you need to define the subnet and captive portal IP for the site. Optionally, you can define an Acceptable Use Policy that overrides the default AUP for the customer. Additional attributes are available–see "tenant add site <attribute>" on page 1380.

   ```
   tenant add site 1 subnet 10.100.1.0 255.255.255.0
   tenant add site 1 redirect 1.2.3.4
   tenant add site 1 aup "aup/site1_aup.htm"
   ```
8. Enable the captive portal feature. See "tenant enable" on page 1382.

   ```
   tenant enable
   ```

Sky UI Tasks Step-by-Step

Step-by-step instructions on how to accomplish a specific task in the Sky UI.

Sky UI Classification Tasks

Before you can manage bandwidth allocation, you must determine exactly what applications are running on your network. PacketShaper s traffic discovery feature can automatically provide you with a list, called a *traffic tree*, of all the applications on a network. PacketShaper differentiates one application from another by evaluating characteristics in traffic flows and grouping them into *classes*. Each traffic class contains at least one *matching rule*, which is a set of characteristics that identifies a specific traffic type. For more information, see the Traffic Classification Overview.

PacketShaper can automatically create classes as traffic is discovered on your network, or you can create classes manually. This section includes tasks for creating traffic classes manually.

Create a Traffic Class in the Sky UI

The traffic classes that PacketShaper auto-discovers are based on service or protocol. If you want to categorize your traffic by IP address, subnet, port, class of service, interface, or VLAN, you can manually create the class in the Sky UI. For example, you might want to create classes for each of your servers so that you can set appropriate QoS policies for them.

The Sky UI allows you to create classes based on multiple criteria, but if you want to create a simple class based on a single matching rule criterion, you can use the "simple match" form. For more complex classes, see Create a Traffic Class with Multiple Criteria.

To create a "simple match" class:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane.

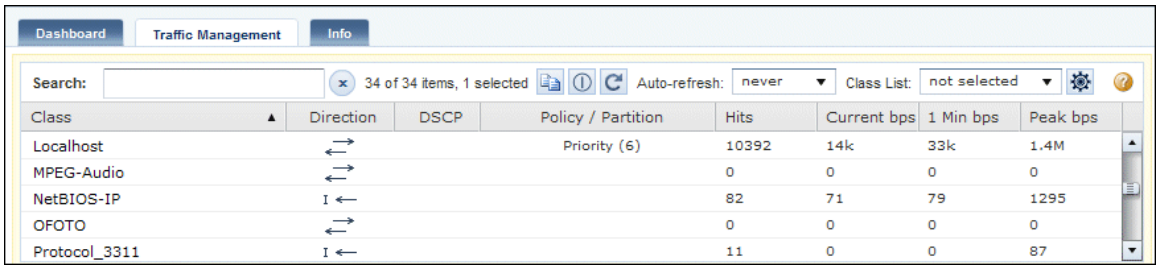

2. In the class tree, select the parent class for your new class (such as Root).

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Add Class**.

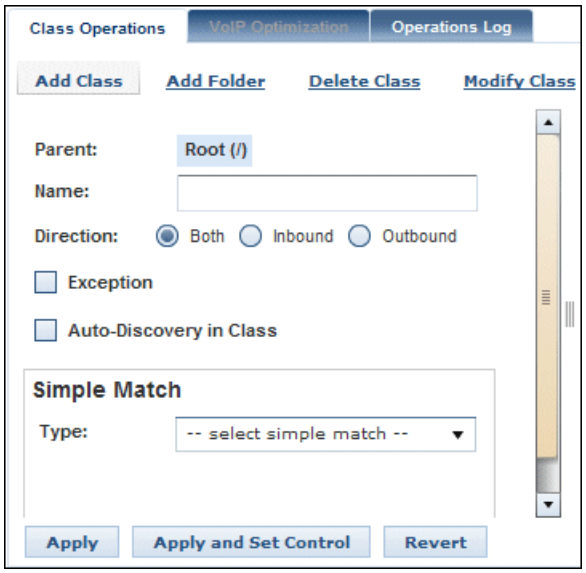

6. In the **Name** field, type a descriptive name for the traffic class (up to 31 characters; hyphens, underscores, and periods are acceptable).

7. Select the **Direction** for the new class: **Both**, **Inbound**, or **Outbound**

8. To create an exception class, enable the **Exception** checkbox.

9. If you want PacketShaper to discover subclasses within the class, enable the **Auto-Discovery in Class** checkbox.

10. From the **Type** drop-down list, select one of the following types of matching rules and then specify the value for the rule:

Type	Value
User Group	Enter all or any part of the group name in the **User Group** field and click the Search icon; then select the name from the results list. Or, to list all group names, leave the field empty and click the icon. See also Create a User Group Class. **Note:** If you manually enter the group name without using the search tool, you must also include the domain name (for example, *CAL\group-engineering*).
User	Enter all or any part of the user name in the **User** field and click the Search icon; then select the name from the results list. Or, to list all user names, leave the field empty and click the icon. See also Create a User-Based Class. **Note:** If you manually enter the user name without using the search tool, you must also include the domain name (for example, *CAL\john.smith*).
URL Category	Select the category name from the drop-down list. See URL Categories for a list of categories, descriptions, and example URLs. See also Create a Category-Based Class.
Operation	Select the web operation from the drop-down list.
Service Group	Select the service group name from the drop-down list. See Applications, Protocols, and Services for a list of service groups and descriptions.
Service	Select the service name from the drop-down list. The list includes both IPv4 and IPv6 services. See Applications, Protocols, and Services for a list of services and descriptions.
DSCP	Select one or more DSCP values (0-63).
IP Address	Enter an Inside and/or Outside IPv4 or IPv6 address, for example, *207.78.98.254 or 2001:db8:1234:5678::2*. Or specify a range of IPv4 addresses by separating the first and last addresses with a dash – for example, *10.10.10.10-10.10.10.20*.
Subnet	Enter the subnet's Inside and/or Outside IPv4 address and CIDR; the CIDR number specifies the number of constant bits in the address range, for example, *192.168.10.0/24*. Or, to specify an IPv6 subnet, enter the IPv6 prefix and prefix length, for example, *2001:db8:2::/48*.
VLAN ID	Enter the VLAN ID (0-4095).
Device	Select an interface (such as a NIC) to classify by.

11. Click **Apply** or **Apply and Set Control**.

Additional Information

- Use the **Revert** button to clear all the values you have entered on the Add Class form (before applying the changes).
- If you get an error message, it's possible that a class already exists with the rule you specified. To get additional help in determining why a class wasn't created, you can look at the operation log.

Exception Classes

When the class tree is in classification order, classes are listed in order from most specific to least specific; classes with more specific criteria are at the top of the tree. Using the class tree's classification order, PacketShaper classifies a traffic flow by searching the class tree, beginning at the top, to find a match for the flow. As soon as PacketShaper finds a match, it categorizes the flow into the appropriate class and looks no further down the tree (even if there may be another class that matches the traffic flow). When you want to change the order that PacketShaper uses to find a match for a traffic flow, you can create *exception classes*.

> **Note:** This discussion assumes the class tree is NOT in combined view. Although you can designate exception classes when the class tree is in combined view, exception classes pertain to classification; therefore, when creating exception classes, it's more useful to display the class tree in classification order than in alphabetical order (combined view). To view the class tree in classification order, see Configuring Class Tree Settings in the Sky UI.

An exception class sorts above standard classes in the class tree. Exception classes give you the ability to redefine the search order that the PacketShaper uses to find a match for a traffic flow. Typically, you do not need to change the position that PacketShaper defines for a class – this is for special exceptions only (hence, the name exception).

Discovery within a Class

Certain classes – such as PVC, VLAN, subnet, Citrix, DICOM, Oracle, PostgreSQL, RTCP, and RTP classes – can be more specifically classified. For these types of classes, you can choose to enable traffic discovery within the class. If you turn on traffic discovery:

- for a subnet class (such as Marketing or Engineering), PacketShaper will discover all the traffic running on the subnet and create child classes for these services.
- for Oracle and PostgreSQL classes, PacketShaper will create child classes for each database that is discovered.
- for Citrix classes, PacketShaper will create child classes by published application.
- for DICOM classes, PacketShaper will automatically discover DICOM client or server applications.
- for RTP and RTCP classes, PacketShaper will automatically create child classes by encoding name (codec).

Traffic discovery cannot be enabled within some classes because the auto-discovery process cannot categorize those traffic types in any greater detail. For example, a class that matches on IPX cannot be further auto-classified.

Create a Traffic Class with Multiple Criteria

To create a class that contains multiple criteria:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

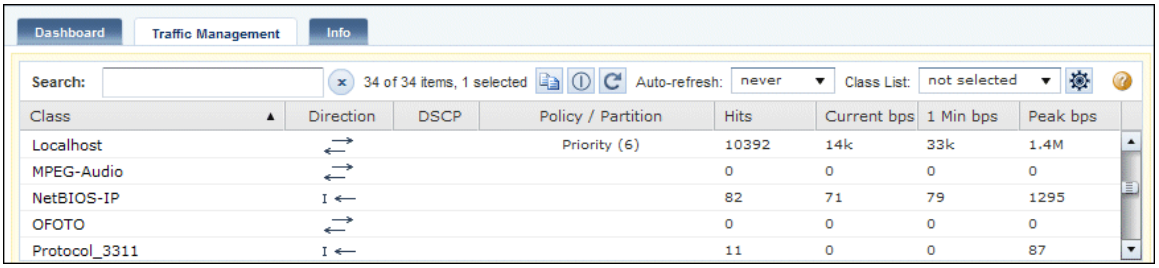

2. In the class tree, select the parent class for your new class (such as Root).

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Add Class**. **show screen...**

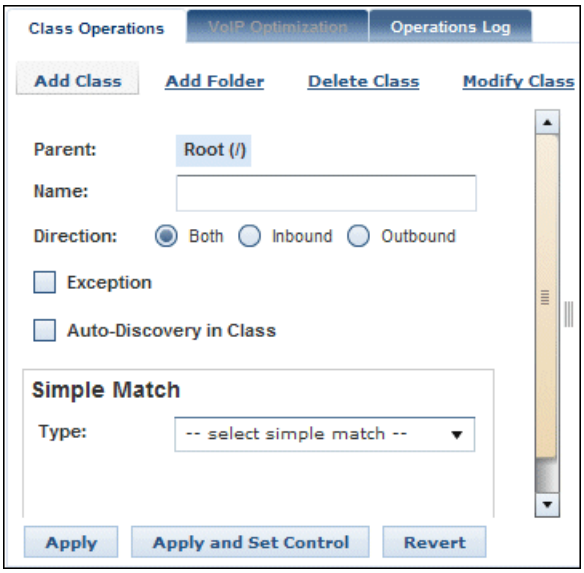

6. In the **Name** field, type a descriptive name for the traffic class (up to 31 characters; hyphens, underscores, and periods are acceptable).

7. Select the **Direction** for the new class: **Both**, **Inbound**, **Outbound**

8. To create an exception class, enable the **Exception** checkbox.

9. If you want PacketShaper to discover subclasses within the class, enable the **Auto-Discovery in Class** checkbox.

10. In the *Matching Rule Criteria* form, select each applicable matching rule criteria and then specify the value for the criteria.
show screen...

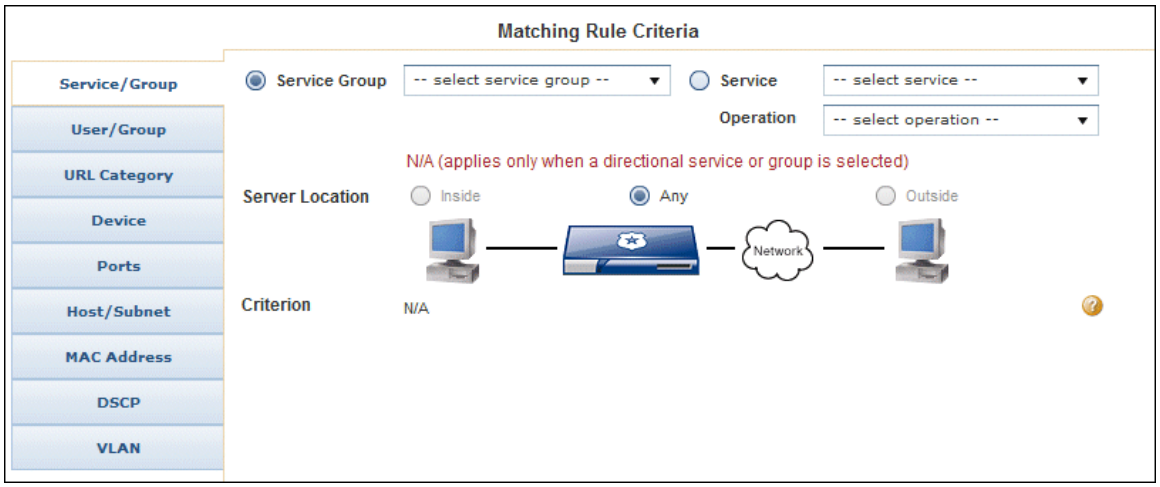

Criteria	Value
Service/Group	**Service:** Select the service name from the drop-down list. The list includes both IPv4 and IPv6 services. *or* **Service Group:** Select the service group name from the drop-down list. The **Criterion** field is available for certain services; this field allows you to further differentiate these types of traffic. For example, web traffic can be differentiated by host DNS name or address, URL, content type, or web browser (user agent). For details, see Application-Specific Matching Rule Criteria. See Applications, Protocols, and Services for service and service group descriptions. **Server Location:** Specify the application's Server Location (**Inside** or **Outside**), relative to the unit. Typically, *inside hosts* are on your LAN and *outside hosts* are on the WAN or Internet. If you wish to capture traffic in a given class regardless of server location, select **Any**. PacketShaper will create two matching rules – one for inside and one for outside. **Tip:** When setting Inside or Outside, answer the question: "If the chosen service uses a server, is the server found on the inside or outside of the PacketShaper unit?"
User/Group	Select **Group**, **User**, or **User List**. For details, see Create a User Group Class, Create a User-Based Class or Create a User List. **Note**: If you manually enter the user or group name without using the search tool, you must also include the domain name (for example, *CAL\john.smith*).
URL Category	Select the category name from the drop-down list. See URL Categories for a list of categories, descriptions, and example URLs. See also Create a Category-Based Class. To create a class that discovers child classes for URL categories, choose **All Categories**, and enable **Auto-Discovery in Class**. See Auto-Discover Category Classes.
Device	Select an interface pair to classify by (such as Slot2).

Criteria	Value
Ports	In the **Inside** and **Outside** fields, list the port(s) associated with the services for this class. In general, you should specify a port number only if you want to restrict classification to a specific port. Since many applications no longer limit transmissions to assigned port numbers, it is better to specify a service name. To specify a range of port numbers—for example, 5001 through 5005 – use a dash (-) to separate the low and high values of the range (5001-5005). To specify non-contiguous port numbers – for example, only ports 5001 and 5025—define separate matching rules within the same class.
Host/Subnet	**Name:** Enter the domain name. (This option requires that you configure a DNS server in PacketShaper.) If you configured a default domain name, you can use non-fully qualified domain names. If the name does not exist when you create the matching rule, PacketShaper warns you about the unknown name. **IP Address:** Enter an Inside and/or Outside IPv4 or IPv6 address, for example, *207.78.98.254*or *2001:db8:1234:5678::2*. Or specify a range of IPv4 addresses by separating the first and last addresses with a dash – for example, *10.10.10.10-10.10.10.20*. **Note:** When specifying an IPv6 address, make sure the service for the class is an IPv6 service. Likewise, when specifying an IPv4 address, the service must be IPv4. **Host List:** Enter the name of a previously-defined host list. (A host list contains IPv4 addresses, DNS names, and/or subnets that traffic class matching rules can reference. For details on creating and using host lists, see Use Host Lists. **Subnet:** Enter the subnet's Inside and/or Outside IP address and CIDR; the CIDR number specifies the number of constant bits in the address range, for example, *192.168.10.0/24*. Or, to specify an IPv6 subnet, enter the IPv6 prefix and prefix length, for example, *2001:db8:2::/48*.
MAC Address	Enter a MAC address to enable classification of traffic to and from a known host. The address must be entered in six-octet format for example, *08:12:34:56:A1:5C*. When the host resides outside PacketShaper, enter the MAC address in the **Outside** field. If the host is on the inside of PacketShaper, use the **Inside** field. If you wish to match all traffic from a host, regardless of whether it's inside or outside of the unit, you must create two separate matching rules. This functionality allows you to create classes for traffic going through a specific router and/or router interface. For example, suppose PacketShaper is connected to two WAN routers. With this feature, you can create a separate class for each router by specifying the router's interface MAC address in the matching rule. You can then assign a partition to each class to control bandwidth usage of the link.
DSCP	Select one or more DSCP values (0-63).
VLAN	Enter the VLAN ID (0-4095) to classify application traffic within multiple virtual LANs on a single ethernet segment (an 802.1q trunk). For a multiple VLAN environment (where packets passing through PacketShaper will have more than one VLAN header), only the outermost VLAN header is used for classification. The other headers are ignored.

11. Click **Apply** or **Apply and Set Control**.

Additional Information

- Use the **Revert** button to clear all the values you have entered on the Add Class form (before applying the changes).
- If you get an error message, it's possible that a class already exists with the rule you specified. To get additional help in determining why a class wasn't created, you can look at the operation log.

Classify Users in the Sky UI

PacketShaper has the ability to monitor and control traffic of Active Directory users and groups. You can create classes based on users or groups for the purposes of monitoring and shaping traffic on the network. This functionality allows administrators to monitor users of a service, monitor what services a specific user or group is using, or manage different policies for a different set of users or groups on the same service.

Create a User-Based Class

Create a User Group Class

Create a Class Based on a User List

Configure PacketShaper as a BCAAA Client

Create a User-Based Class in the Sky UI

When the Blue Coat Authentication and Authorization Agent (BCAAA) is installed on a Windows Domain Controller and the PacketShaper is configured to be a BCAAA client, you can create traffic classes for any of your users in Active Directory. Once you have created a user class, you can apply controls to that user's traffic.

To create a class for a single user:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane.

2. In the class tree, select the parent class for your new class (such as Root).

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Add Class**.

6. In the **Name** field, type a name to identify the user (up to 31 alphanumeric characters; hyphens, underscores, and periods are acceptable; spaces and UTF-8 characters are not allowed).

7. For **Direction**, select **Both**.

8. For the Simple Match Type, select **User**.

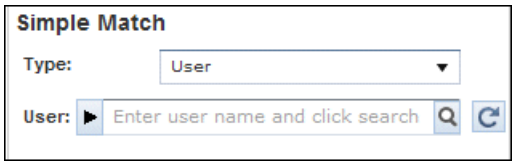

Or, in the Matching Rule Criteria form, you can display the **User/Group** panel and select **User**.

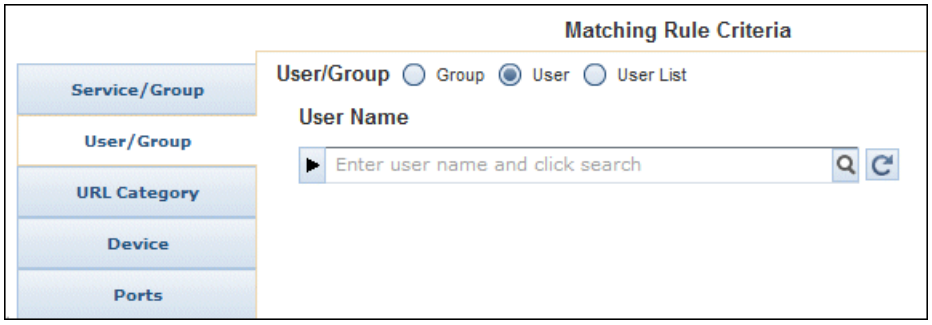

9. Specify the user name using one of the following methods:

- Enter all or any part of the user name (such as the first or the last name) in the **User** field. For example, enter *john*, *smith*, or *john.smith*. Then click the Search icon and select the name from the search results list. Note that you may enter UTF-8 characters for the user name (but not for the class name).

- To select the user name from a list all users, leave the **User** field empty and click the icon; then select the name from the search results list.
- To list all user names for a specific domain, enter the domain name and click the icon; then select the name from the search results list.

> **Note:** If the search results list is empty, no users were found that match the criteria you entered. Double-check your typing and spelling and try again. Experiment with other types of searches, such as searching by just the first name or just the last name, or display the list of all users and pick the user from the list.

10. Click **Apply**.

Additional Information

- When creating classes, PacketShaper does not validate that the specified user or group exists in Active Directory. In addition, PacketShaper allows you to create a child user-based class of a parent group-based class, even if the user isn't in the AD group of the parent class. If you create a class for a user or group that doesn't exist in AD, the class will not have a configuration error, but no traffic will hit the class.
- You can sort the search results by user or domain name in either ascending or descending order; to sort, click the User or Domain column headings.
- To clear the **User** field of a search string you previously entered, click the button.
- Use the disclosure triangles, and , to expand and collapse the results list.
- When you display the list of all users, these user names are cached locally for the current browser session. To force a refresh of the list (in case users are added, deleted, or modified during your session), click the Refresh button. You will also need to refresh if you change the BCAAA configuration during your current browser session.

Create a User Group Class in the Sky UI

When the Blue Coat Authentication and Authorization Agent (BCAAA) is installed on a Windows Domain Controller and the PacketShaper is configured to be a BCAAA client, you can create traffic classes for any of your user groups in Active Directory. Once you have created a user group class, you can apply controls to that group's traffic.

To create a class for a use group:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane.

2. In the class tree, select the parent class for your new class (such as Root).

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Add Class**.

6. In the **Name** field, type a name to identify the user group(up to 31 alphanumeric characters; hyphens, underscores, and periods are acceptable; spaces and UTF-8 characters are not allowed).

7. For **Direction**, select **Both**.

8. For the Simple Match Type, select **User Group**.

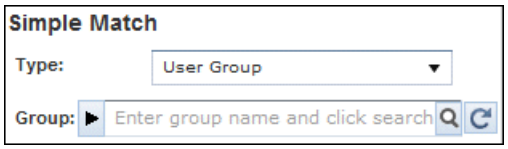

Or, in the Matching Rule Criteria form, you can display the **User/Group** panel and select **Group**.

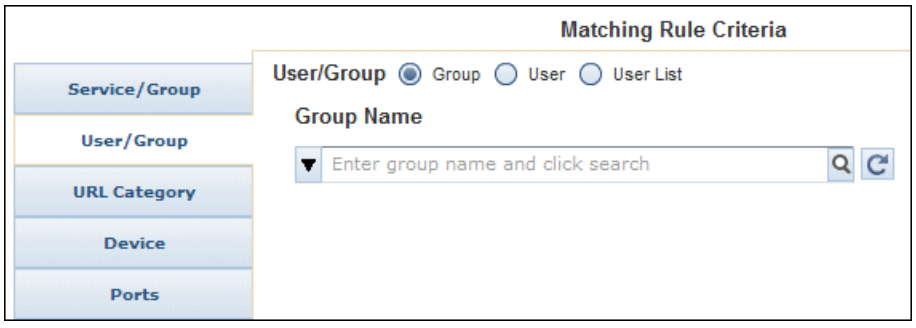

9. Specify the group name using one of the following methods:

- Enter all or any part of the user group name in the **Group Name** field. For example, enter *group-sales*, or *sales*. Then click the Search icon and select the name from the search results list. Note that you may enter UTF-8 characters for the group name (but not for the class name).

- To select the user group name from a list all groups, leave the **Group Name** field empty and click the icon; then select the name from the search results list.
- To list all group names for a specific domain, enter the domain name and click the icon; then select the name from the search results list.

> **Note:** If the search results list is empty, no groups were found that match the criteria you entered. Double-check your typing and spelling and try again. Experiment with other types of searches, or display the list of all groups and pick the group from the list.

10. Click **Apply**.

Additional Information

- When creating classes, PacketShaper does not validate that the specified user or group exists in Active Directory. In addition, PacketShaper allows you to create a child user-based class of a parent group-based class, even if the user isn't in the AD group of the parent class. If you create a class for a user or group that doesn't exist in AD, the class will not have a configuration error, but no traffic will hit the class.
- You can sort the search results by group or domain name in either ascending or descending order; to sort, click the Group or Domain column headings.
- To clear the **Group Name** field of a search string you previously entered, click the button.
- Use the disclosure triangles, and , to expand and collapse the results list.
- When you display the list of all groups, these group names are cached locally for the current browser session. To force a refresh of the list (in case groups are added, deleted, or modified during your session), click the Refresh button. You will also need to refresh if you change the BCAAA configuration during your current browser session.

Create a Class Based on a User List in the Sky UI

When the Blue Coat Authentication and Authorization Agent (BCAAA) is installed on a Windows Domain Controller and the PacketShaper is configured to be a BCAAA client, you can create a named list of users and groups, then create a traffic class based on this user list. This allows you to monitor and control traffic for a set of users or groups. (For creating a class based on a single user or group, see Create a User-Based Class or Create a User Group Class.)

To create a class for a user list:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane.

2. In the class tree, select the parent class for your new class (such as Root).

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Add Class**.

6. In the **Name** field, type a name to identify the user list (up to 31 alphanumeric characters; hyphens, underscores, and periods are acceptable; spaces and UTF-8 characters are not allowed).

7. For **Direction**, select **Both**.

8. In the Matching Rule Criteria form, display the **User/Group** panel.

9. Select **User List**. The User List Editor displays.

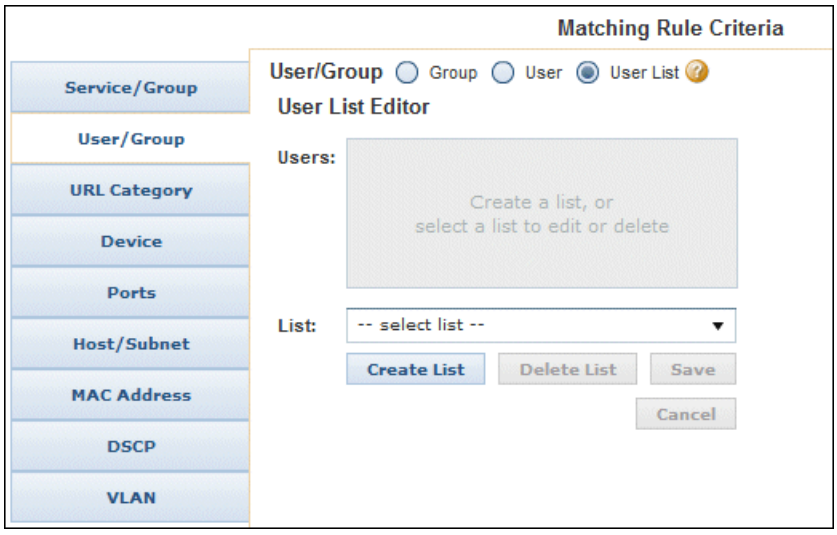

10. Click **Create List** to create a new user list, or select an existing one from the **List** drop-down.

11. Click **Apply**.

Open the User List Editor in the Sky UI

The Sky UI offers a User List Editor for creating and managing user lists. There are two ways to open the User List Editor:

- In the Policy Manager, without having a specific class selected
- While creating or modifying a class or a matching rule

To open the standalone User List Editor:

1. Click the **Traffic Management** tab.

2. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

3. In the **Class Operations** tab, click **User List Editor**. The User List Editor appears in the Policy Manager pane.

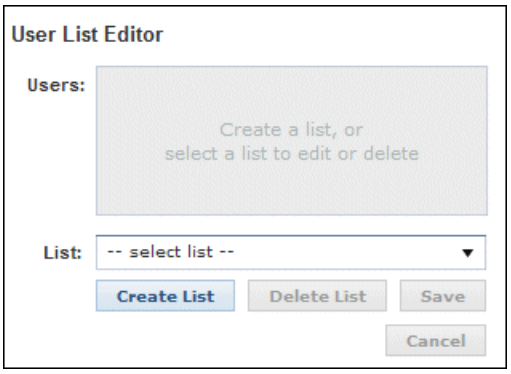

4. You can now create, modify, or delete user lists.

To open the User List Editor while creating/modifying a class or a matching rule:

1. While creating/modifying a class or a matching rule, display the **User/Group** panel (in the Matching Rule Criteria section).

2. Click **User List**. The User List Editor appears.

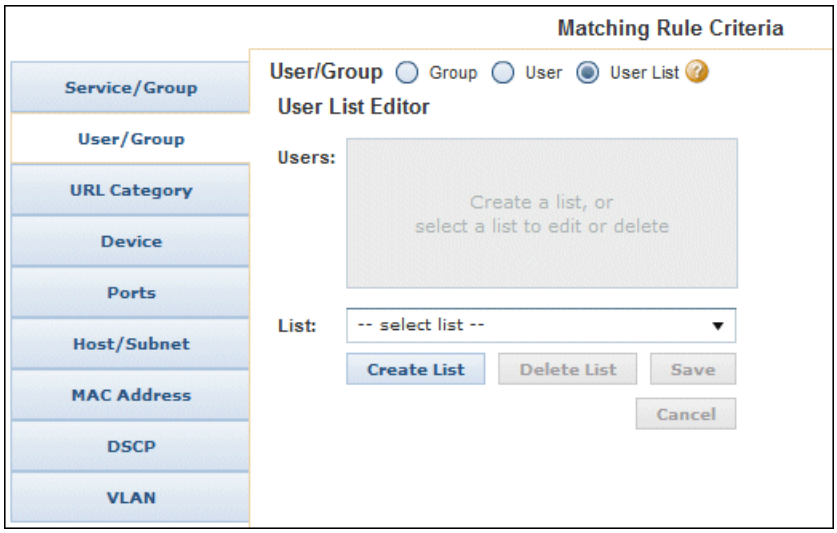

3. You can now create, modify, or delete user lists.

Create a User List in the Sky UI

The Policy Manager in the Sky UI offers a User List Editor for creating user lists. You can either enter the user and group names manually, paste them in from another source, or select them from a list.

To create a user list:

1. Open the User List Editor.

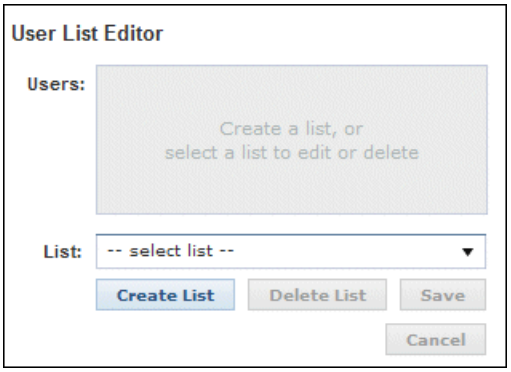

2. Click **Create List**.

3. If you know the exact user or group names, you can type them in the **Users** text box; separate each name with a comma, a tab, or a new line. You can also paste the names into the text box. Precede each user name with **u:** and each group name with **g:**. Note that you must also type the domain name. Examples: *u:cal\joan.thomas g:ny\group-finance*.

4. If you want to select the names from a list, click **Add Users**. The list of user names appears on the right.

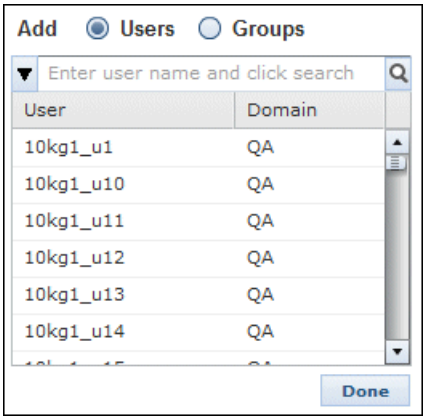

5. Select **Users** or **Groups**, then use any of the following methods to locate names:

- Use the scroll bar to locate the desired names on the search results list. To select multiple names, use Ctrl+click or Shift+click

- Enter all or any part of the user or group name (such as the first or the last name) in the **Add** field. For example, enter *john*, *smith*, or *john.smith* for a user name, or *sales or group-sales*for a group name. Then click the Search icon and select the name from the search results list. Note that you may enter UTF-8 characters for the user or group name.
- To list all user or group names for a specific domain, enter the domain name and click the icon; then select the names from the search results list.

Note: If the search results list is empty, no users or groups were found that match the criteria you entered. Double-check your typing and spelling and try again. Experiment with other types of searches, or display the list of all users/groups and pick the name from the list.

6. Click **<-Add**. The names are added to the **Users** box on the left. Note that the system automatically added **u:** before each user name and **g:** before each group name.

7. In the **List Name** field, enter a descriptive name for the list and click **Save**.

You can now specify this user list when creating classes or adding matching rules.

Additional Information

- The User List Editor does not validate the existence of the names.
- A Sky user list can contain up to 612 characters. If you have more users than can fit in a single list, create multiple user lists and add a matching rule for each one.

Delete a User List in the Sky UI

You can use the User List Editor to delete a user list that is no longer needed. Note that the Editor prevents you from deleting a user list that is currently specified in a matching rule.

To delete a user list:

1. Open the User List Editor.

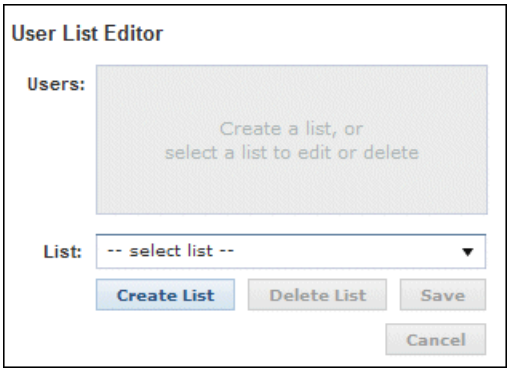

2. In the **List** field, select the user list to delete. The user and group names associated with this list are displayed in the **Users** text box.

3. Click **Delete List** and click OK to confirm.

> **Note:** If you try to delete a user list that is currently specified in a matching rule, you will see the following message: *User list xxxx is in use*. Before you can delete this user list, you need to determine which classes specify this user list in a matching rule.

Edit a User List in the Sky UI

You can use the User List Editor to add, modify, and delete user or group names after you have created a user list. A user list created in Sky can contain up to 612 characters. If you have more users/groups than can fit in a single list, create multiple user lists and add a matching rule for each one.

To modify a user list:

1. Open the User List Editor.

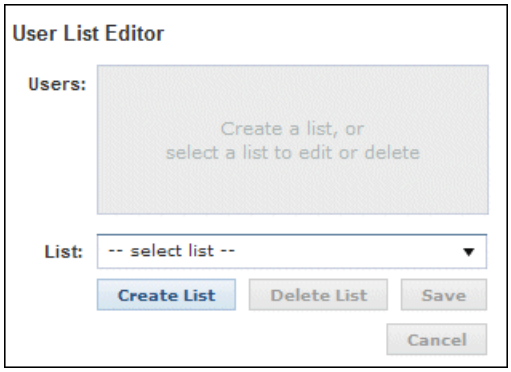

2. In the **List** field, select the user list to modify. The user and group names associated with this list are displayed in the **Users** text box.

> **Note:** If you are opening a user list that was created in the CLI or Advanced UI, it may have more characters than the maximum allowed in Sky (612). You will see a warning message if this happens; you should not edit these lists in Sky because the changes cannot be saved.

3. To add another user or group name to this list, enter it in the **Users** text box. Separate each name with a comma, a tab, or a new line. Precede each user name with **u:** and each group name with **g:**. Note that you must also type the domain name. Examples: *u:cal\joan.thomas g:ny\group-finance*. You can also paste the user and group names into the text box.

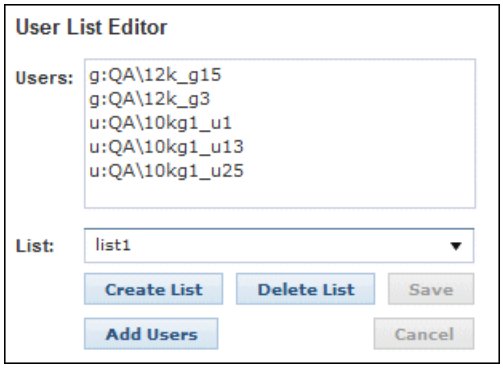

4. If you want to select the names from a list, click **Add Users**. The list of user names displays on the right.

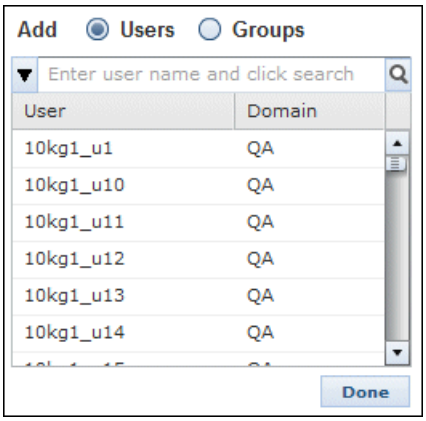

5. Select **Users** or **Groups**, then use any of the following methods to select names:

- Use the scroll bar to locate the desired names on the search results list. To select multiple names, use Ctrl+click or Shift+click,
- Enter all or any part of the user or group name (such as the first or the last name).For example, enter *john*, *smith*, or *john.smith* for a user name, or *sales* or *group-sales* for a group name. Then click the Search icon and select the name from the search results list. Note that you may enter UTF-8 characters for the user or group name.
- To list all user or group names for a specific domain, enter the domain name and click the icon; then select the names from the search results list.

> **Note:** If the search results list is empty, no users or groups were found that match the criteria you entered. Double-check your typing and spelling and try again. Experiment with other types of searches, or display the list of all users/groups and pick the name from the list.

6. Click **<- Add**. The names are added to the **Users** box on the left. Note that the system automatically added **u:** before each user name and **g:** before each group name.

7. To modify a user or group name, you can edit the text directly in the **Users** box. For example, you can fix a typing mistake by inserting and deleting letters.

8. To remove a user or group name, select the text in the **Users** box and press delete.

9. Click **Save** to save your changes.

> **Note:** The User List Editor does not validate the existence of the names.

Manually Create an IPv6 Class in the Sky UI

To create an IPv6 class in the Sky UI, just choose the appropriate IPv6 service (such as HTTP-IPv6 or FTP-IPv6).

1. Click the **Traffic Management** tab. The class tree appears in the upper pane.

2. In the class tree, select the parent class for your new class (such as Root).

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Add Class**.

6. In the **Name** field, type a descriptive name for the traffic class (up to 31 characters; hyphens, underscores, and periods are acceptable).

7. Select the **Direction** for the new class: **Both**, **Inbound**, or **Outbound**

8. For the **Service/Group** matching rule criteria, select **Service** and choose the IPv6 service name from the list. IPv6 services have *IPv6* appended to the name.

9. Click **Apply**.

Classify URL Categories in the Sky UI

PacketShaper offers a number of ways you can classify by URL categories:

Create a Category-Based Class

Auto-Discover Category Classes

Discover Categories Under Web-Based Classes

Create a Class for Multiple Categories

Create a Category-Based Class in the Sky UI

You can manually create classes for any of the 70+ categories PacketShaper can classify. Alternatively, you can have PacketShaper auto-discover the categories.

To create a category-based class:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane.

2. In the class tree, select the parent class for your new class (such as Root).

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Add Class**.

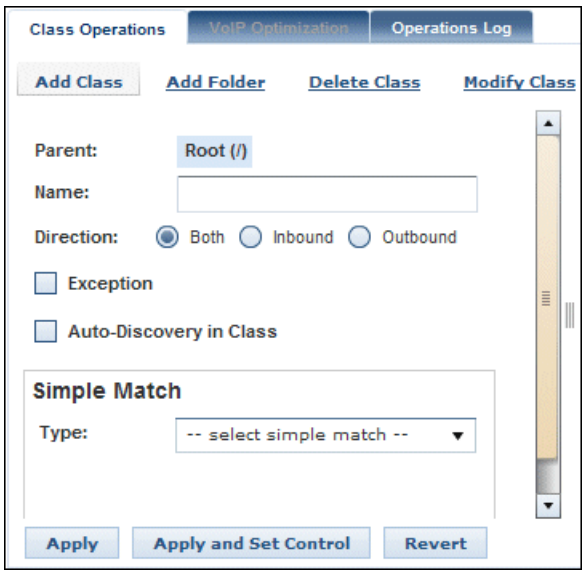

6. In the **Name** field, type a name for the category class (up to 31 characters; hyphens, underscores, and periods are acceptable).

7. For **Direction**, select **Both**.

8. (Optional) Select **Exception**.

> **Note:** When a category class is created in the root, you will usually want to make it an exception class to make sure the web traffic gets classified into the category class instead of in service-based classes (such as HTTP, Facebook, MySpace, or other services that use

> HTTP as a transport).

9. For the Simple Match Type, select **URL Category**, and select the category name from the drop-down list.

10. Click **Apply**.

11. Make sure WebPulse is enabled.

Auto-Discover Category Classes in the Sky UI

To auto-discover category classes, you need to create an "all categories" class, turn on class discovery, and make sure global traffic discovery is enabled. If you also want to auto-discover service-based classes, you should make the "all categories" class an exception so that it appears at the top of the traffic tree and web traffic will get classified into the category class instead of the web-based classes (such as HTTP, SSL, or HTTP-Tunnel).

To auto-discover category classes:

1. In the Root(/), create a simple-match class named Categories.

2. For **Direction**, select **Both**.

3. Select **Exception**.

4. Enable **Auto-Discovery in Class**.

5. For the Simple Match Type, select **URL Category**, and select **All Categories** from the drop-down list.

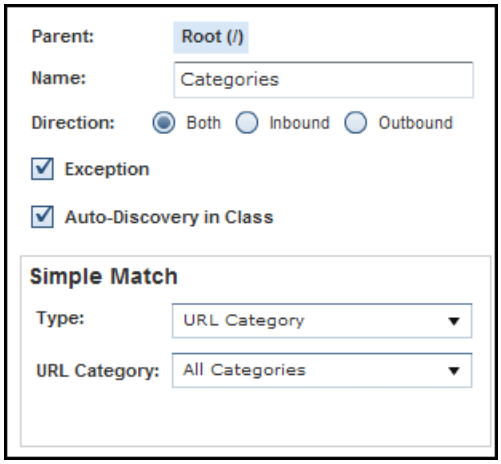

6. Click **Apply**.

7. Make sure WebPulse and global traffic discovery are enabled.

As web traffic appears on your network, PacketShaper will create category classes under the Categories class. For example:

```
Categories
 Business_Economy
 Chat_Instant_Messaging
 Computers_Internet
 Content_Servers
 Email
 Financial_Services
 Shopping
```

```
Social_Networking
Default
```

> **Note:** By default, PacketShaper auto-discovers category classes after it sees one flow for the category in a one-minute time interval. This threshold is user configurable via the **URL Categories** system variable. See Adjust System Variables in PacketGuide.

Discover Categories Under Web-Based Classes in the Sky UI

As an alternative to creating an "all categories" class to auto-discover category classes, you can turn on class discovery for web-based classes (such as HTTP, SSL, HTTP-Tunnel). This technique allows you to have different controls for categories depending on the type of web service. For example, you can create different policies for categories that are secure web traffic versus unsecure.

Although you can auto-discover categories for any web-based service that uses HTTP or HTTPS as transport, Symantec recommends that you limit this feature to HTTP, SSL, and HTTP-Tunnel.

To discover categories under web-based classes:

1. If your class tree doesn't already have classes for HTTP, SSL, and HTTP-Tunnel, create classes for these services.

2. Select **Exception** and enable **Auto-Discovery in Class** for these classes.

3. Make sure WebPulse and global traffic discovery are enabled.

As web traffic appears on your network, PacketShaper will create category classes under HTTP, SSL, and HTTP-Tunnel. For example:

```
HTTP
 Business_Economy
 Chat_Instant_Messaging
 Computers_Internet
 Content_Servers
 Email
 Financial_Services
 Shopping
 Social_Networking
 Default

SSL
 Brokerage_Trading
 Content_Servers
 Default
```

Additional Information

- For category classes to be discovered under an HTTP class, the HTTP class *cannot* have any criteria (such as host, URL, content type, or user agent) defined.
- By default, PacketShaper auto-discovers category classes after it sees one flow for the category in a one-minute time interval. This threshold is user configurable via the **URL Categories** system variable. See Adjust System Variables in PacketGuide.

Create a Class for Multiple Categories in the Sky UI

If you want to apply the same controls to multiple categories (for instance, apply a never-admit policy to all categories containing adult related content), you can create a single class (Adult_Content) and add matching rules for each category. This technique is also useful to consolidate classes in situations where your PacketShaper is reaching its maximum class limit.

To create a class that matches web traffic from multiple categories:

1. In the Root of the class tree, create a class that describes the type of traffic (for example, Adult_Content), and choose the first **URL Category** (such as Pornography) from the drop-down list.

2. Select **Exception**, and click **Apply**.

> **Note:** When a category class is created in the root, you will usually want to make it an exception class to make sure the web traffic gets classified into the category class instead of in service-based classes (such as HTTP, Facebook, MySpace, or other services that use HTTP as a transport).

3. Add a URL category matching rule for each category associated with the class.

4. Make sure WebPulse is enabled.

Create a Web Application Class in the Sky UI

By using the WebPulse cloud service, PacketShaper can expand its classification capabilities to include over 100 web applications, such as Facebook, Gmail, LinkedIn, Twitter, and Flickr. In addition to classifying by web application, you can classify by specific operations (such as Download Attachment, Play Video, and Send Email). Thus, you can create the following types of classes:

- A specific web application for any operation (Service=Facebook, Operation=any)
- A specific operation for any web application (Operation=Upload Pictures, Service=any)
- A specific web application for a specific operation (Service=Facebook, Operation=Upload Pictures)

To create a web application class:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane.

2. In the class tree, select the parent class for your new class (such as Root).

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Add Class**.

6. In the **Name** field, type a name for the web application class (up to 31 characters; hyphens, underscores, and periods are acceptable).

7. In the Matching Rule Criteria form, select **Service/Group**.

8. To classify by a particular web application, choose the **Service** from the drop-down list.

9. *(Optional)* To classify by a specific web operation (such as Download Attachment), choose the **Operation** from the drop-down list. If you selected a web application for the Service in the previous step, the **Operation** drop-down will list only the operations that are applicable to that application.

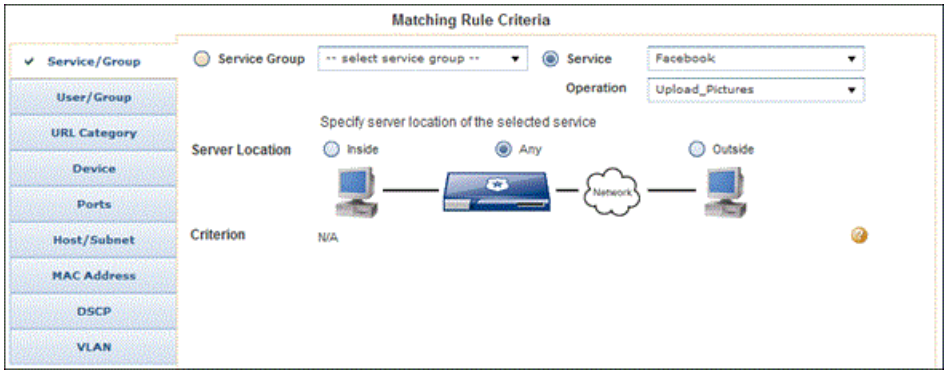

10. Click **Apply**.

11. Make sure WebPulse is enabled.

Create a VoIP Class in the Sky UI

A Voice over IP (VoIP) class is based on the RTP-I service and has a codec assigned to it.

To create a VoIP class:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

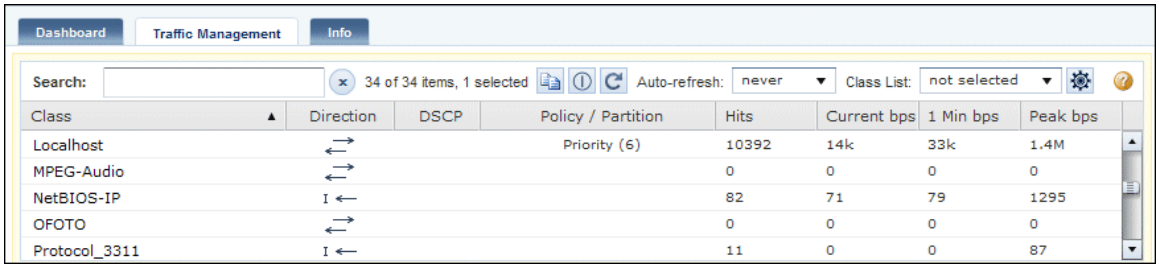

2. In the class tree, select the parent class for your new class (such as Root).

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Add Class**. **show screen...**

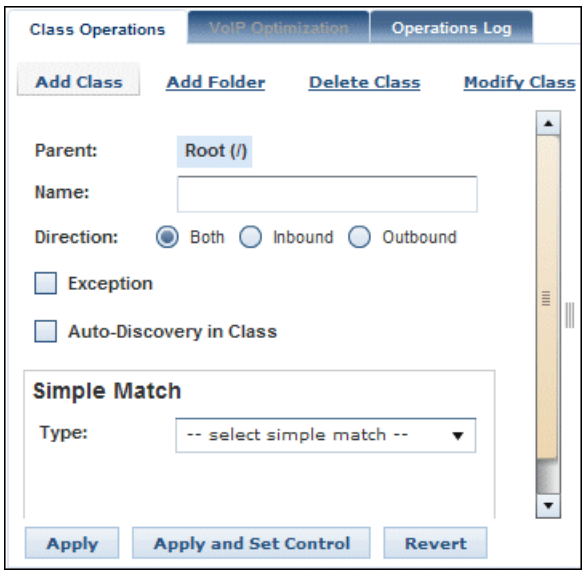

6. In the **Name** field, type a descriptive name for the traffic class (up to 31 characters; hyphens, underscores, and periods are acceptable).

7. Select the **Direction** for the new class: **Both**, **Inbound**, **Outbound**

8. In the *Matching Rule Criteria* form, select the **Service** button.

9. In the **Service** drop-down list, select **RTP-I**. **show screen...**

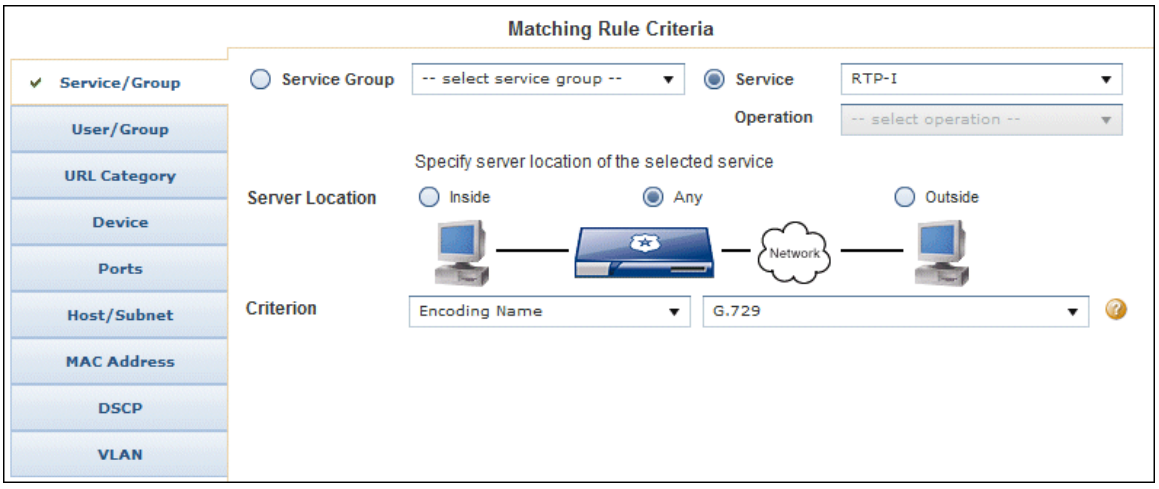

10. In the **Criterion** drop-down list, select **Encoding Name** and then select the codec.

11. Click **Apply.**

You can now set a policy and partition on this class to optimize VoIP calls. See VoIP Optimization.

Modify a Traffic Class in the Sky UI

To modify the settings of an existing class:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

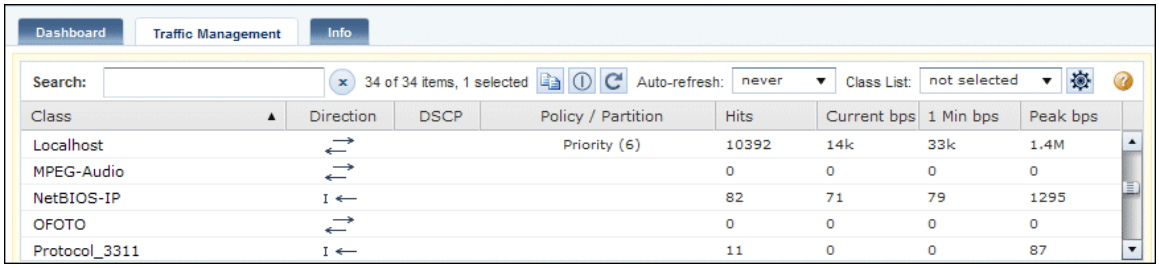

2. In the class tree, select the class you want to modify.

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **Class Operations** tab.

5. Click **Modify Class**. **show screen...**

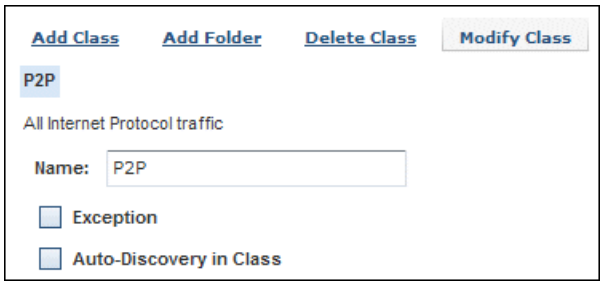

6. To create an exception class, enable the **Exception** checkbox.

7. If you want PacketShaper to discover subclasses within the class, enable the **Auto-Discovery in Class** checkbox.

8. If desired, edit the matching rules.

9. Click **Apply**.

Add Matching Rules in the Sky UI

Matching rules define the criteria the PacketShaper uses to identify traffic types. Every traffic class must have at least one matching rule. When traffic discovery creates a traffic class, it creates one or more matching rules to characterize the traffic type. In a similar fashion, when you create a class manually, you must specify the matching rules that describe the application's flows.

A traffic class can have multiple matching rules, which are treated as separate, distinct specifications. When PacketShaper tries to map a traffic flow to a class, it compares the flow with the criteria in the class' first matching rule. If PacketShaper does not find a match, it continues through the rules until a match is found or until it runs out of matching rules, in which case it moves on to the next class in the tree. If a specific traffic class cannot be found for a flow, the traffic is classified in the Default traffic class for the subtree.

To edit a matching rule:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

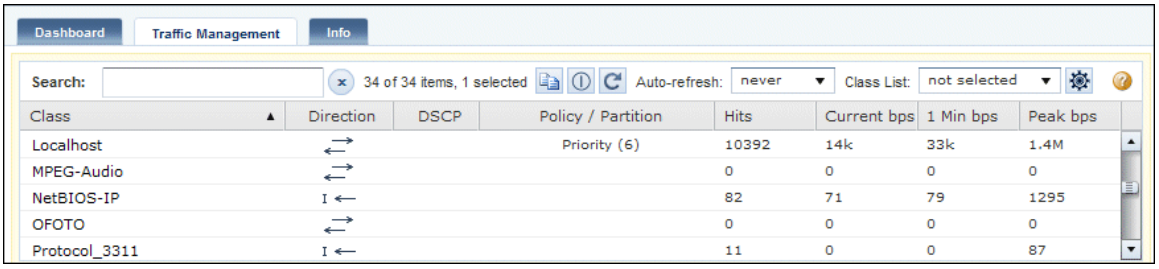

2. Select the class for which you want to edit a matching rule.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Modify Class**.

6. In the *Matching Rules* form, click **Add Rule**. the *Matching Rule Criteria* form displays. **show screen...**

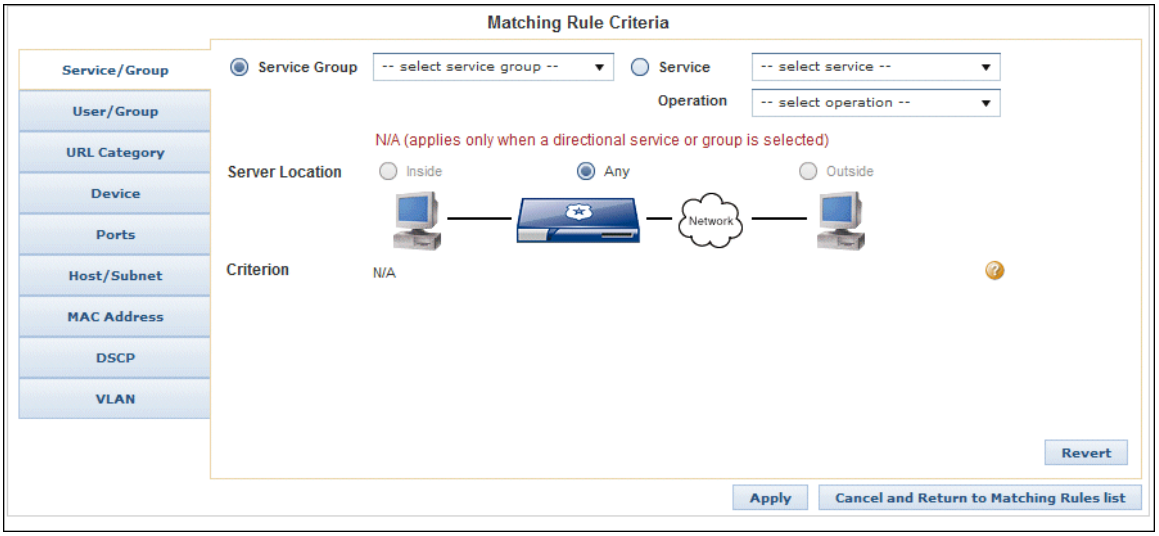

7. In the *Matching Rule Criteria* form, specify criteria as desired. (See Create a Class with Multiple Criteria for descriptions of each criterion.)

8. Click **Apply**.

Additional Information

- To remove a matching rule, select the rule on the *Matching Rules* form and click **Delete Rule**.
- If you get an error message, you can look at the operation log.

Edit Matching Rules in the Sky UI

Matching rules define the criteria the PacketShaper uses to identify traffic types.

To edit a matching rule:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

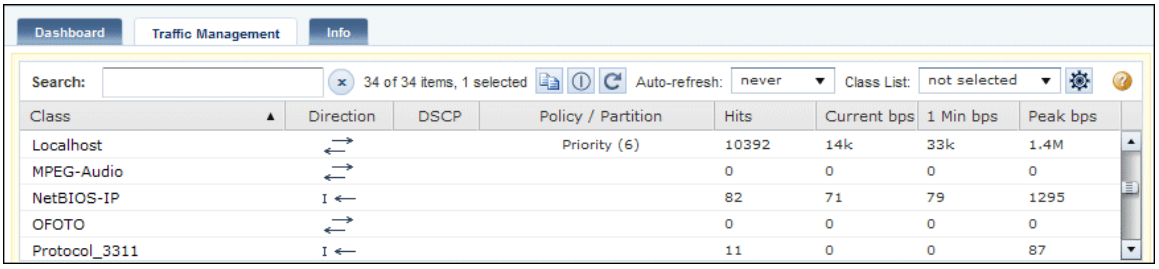

2. Select the class for which you want to edit a matching rule.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Modify Class**. The *Matching Rules* form appears on the right.

6. In the *Matching Rules* form, select the rule you want to modify, and click **Edit Rule**. The *Matching Rule Criteria* form appears. **show screen...**

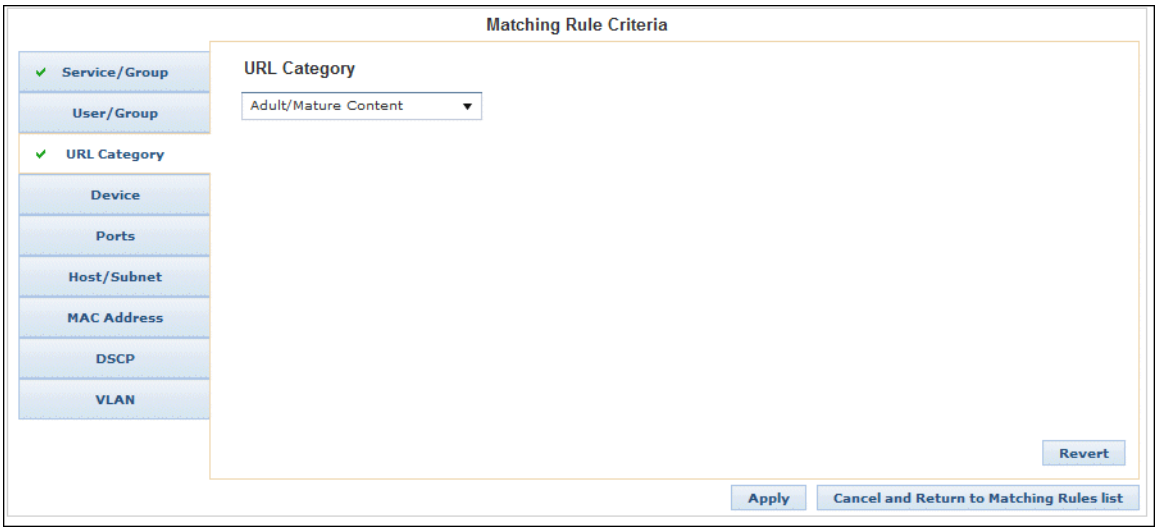

7. Add/modify criteria as desired. (See Creating a Class with Multiple Criteria for descriptions of each criterion.)

8. Click **Apply**.

Additional Information

- To remove a matching rule, select it and click **Delete Rule**.
- If you get an error message, you can look at the operation log.

Use Host Lists in the Sky UI

A *host list* contains a set of IPv4 addresses, DNS names, and/or subnets that traffic class matching rules can reference. Host lists let you specify up to 2048 IPv4 addresses and/or names in a single traffic class matching rule. The addresses in a host list do not need to be contiguous.

The Host List Editor is a tool for creating, updating, and deleting host lists.

Open the Host List Editor

Create a Host List

Edit a Host List

Delete a Host List

Open the Host List Editor in the Sky UI

The Sky UI offers a Host List Editor for creating and managing host lists. There are two ways to open the Host List Editor:

- In the Policy Manager, without having a specific class selected
- While creating or modifying a class or a matching rule

To open the standalone Host List Editor:

1. Click the **Traffic Management** tab.
2. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.
3. In the **Class Operations** tab, click **Host List Editor**. The Host List Editor appears in the Policy Manager pane. **show screen...**

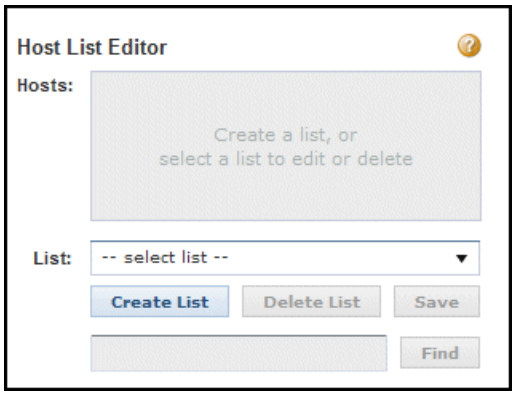

4. You can now create, modify, or delete host lists.

To open the Host List Editor while creating/modifying a class or a matching rule:

1. While creating/modifying a class or a matching rule, click **Host/Subnet** in the Matching Rule Criteria section. **show screen...**

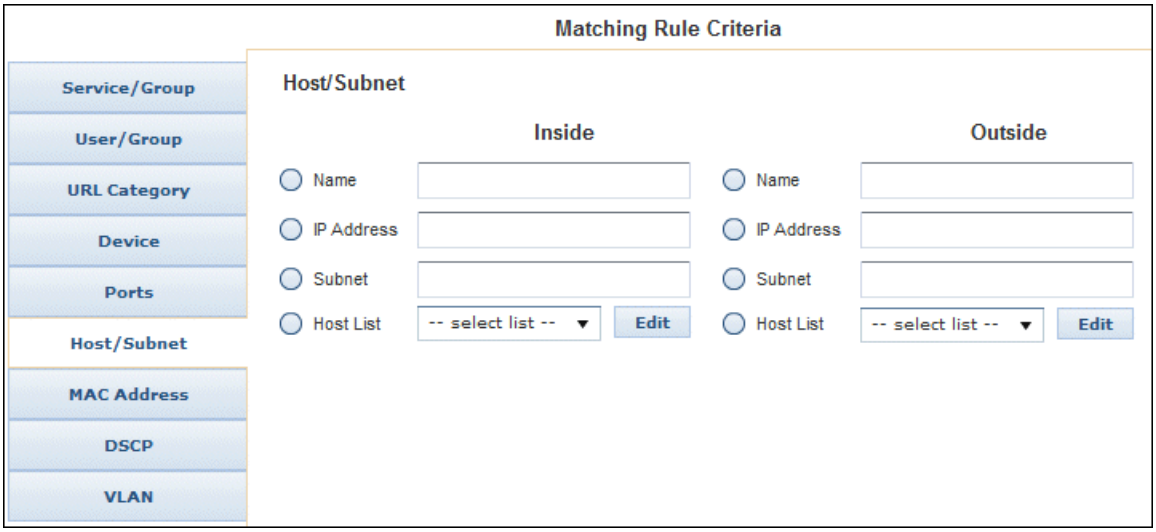

2. Click **Edit**. The Host List Editor opens. **show screen...**

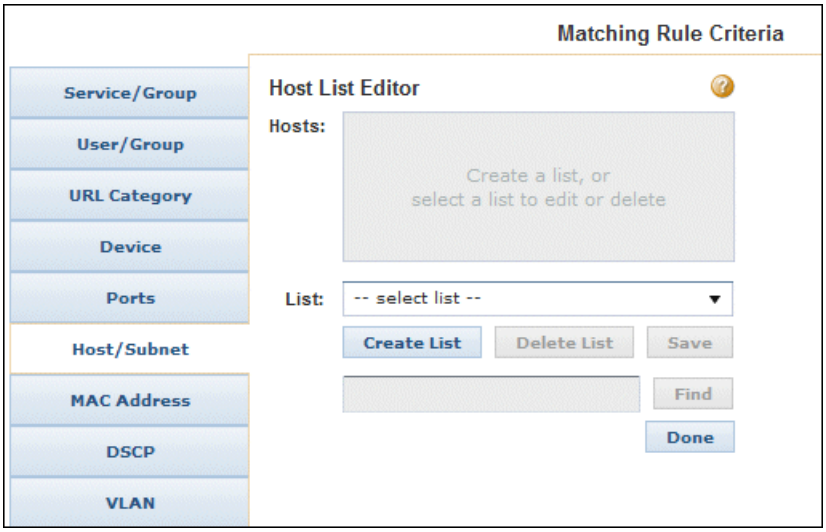

> **Note:** There are two **Edit** buttons: one for Inside and another for Outside. Choose the one appropriate for the side your hosts are on, relative to the PacketShaper.

3. You can now create, modify, or delete host lists.

Create a Host List in the Sky UI

A *host list* contains a set of IPv4 addresses, DNS names, and/or subnets that traffic class matching rules can reference. Host lists let you specify up to 2048 IPv4 addresses and/or names in a single traffic class matching rule. The addresses in a host list do not need to be contiguous.

The Sky UI offers a Host List Editor for creating and managing host lists.

To create a host list:

1. Open the Host List Editor. **show screen...**

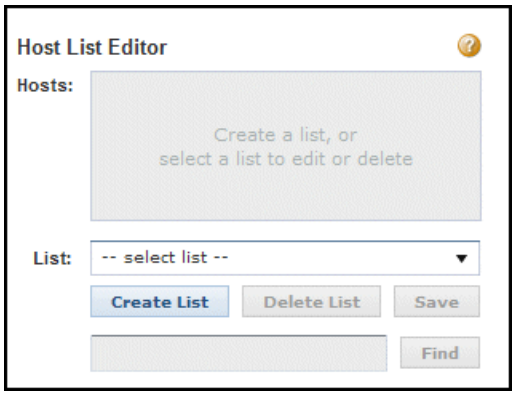

2. Click **Create List**. **show screen...**

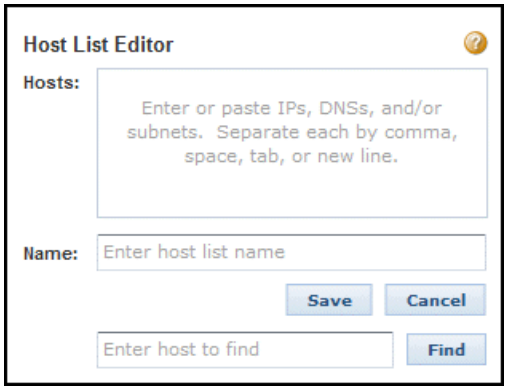

3. In the **Hosts** text box, enter IP addresses, subnets, ranges of IP or subnet addresses, or DNS names; separate each host with a comma, a space, a tab, or a new line. You can also paste the hosts into the text box. Hosts can be specified in any of the following ways:

Type of Host	Example
Host IPv4 address	192.168.1.10
Range of IPv4 addresses Use a dash – with no spaces – between the low and high address in the range.	192.168.1.100- 192.168.1.200

Type of Host	Example
Address of the subnet; the CIDR number specifies the number of constant bits in the address range	192.168.10.0/24
Range of subnet addresses; the CIDR number specifies the number of constant bits in the address range Use a dash between the low and high address in the range. Spaces are not allowed before or after the dash or slash characters.	192.168.10.0- 192.168.20.0/24
DNS name **Note:** Do not use domain names if you will be using the host list with the host sidedness feature.	www.yourcompany.com

4. In the **Name** field, enter a descriptive name for the list and click **Save**.

You can now specify this host list when creating classes or adding matching rules.

Additional Information

- The syntax of your host entries are not validated until you try to save the host list. If you entered an invalid host specification, the Host List Editor indicates which host is invalid so that you can correct the problem and try saving again. If there is an issue with multiple hosts, the Editor alerts you to them one at a time.

- The Host List Editor does not validate the existence of the hosts.

Edit a Host List in the Sky UI

You can use the Host List Editor to add, modify, and delete hosts after you have created a host list. Host lists let you specify up to 2048 IPv4 addresses and/or names in a single traffic class matching rule. The addresses in a host list do not need to be contiguous.

To modify a host list:

1. Open the Host List Editor.

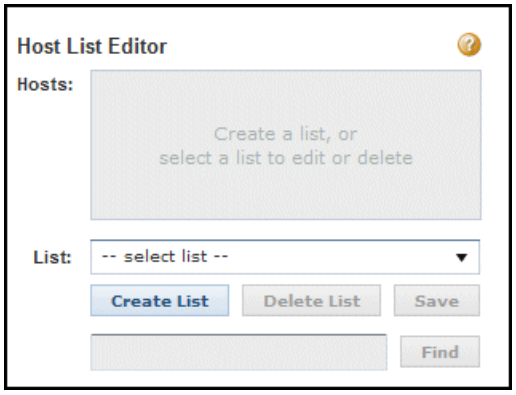

2. In the **List** field, select the host list to modify. The hosts associated with this list are displayed in the **Hosts** text box.

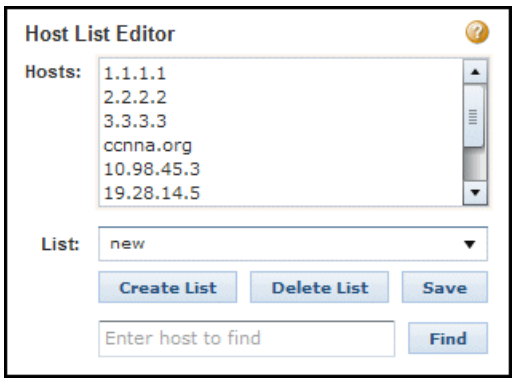

3. To add another host to this list, enter it in the **Hosts** text box. You may enter IP addresses, subnets, ranges of IP or subnet addresses, or DNS names; separate each host with a comma, a space, a tab, or a new line. You can also paste the host specifications into the text box. Hosts can be specified in any of the following ways:

Type of Host	Example
Host IPv4 address	192.168.1.10
Range of IPv4 addresses Use a dash – with no spaces – between the low and high address in the range.	192.168.1.100-192.168.1.200
Address of the subnet; the CIDR number specifies the number of constant bits in the address range	192.168.10.0/24

Type of Host	Example
Range of subnet addresses; the CIDR number specifies the number of constant bits in the address range Use a dash between the low and high address in the range. Spaces are not allowed before or after the dash or slash characters.	192.168.10.0-192.168.20.0/24
DNS name	www.yourcompany.com

Note: Do not use domain names if you will be using the host list with the host sidedness feature.

4. To modify a host specification, you can edit the text directly. For example, you can fix a typing mistake in an IP address by inserting and deleting values.

> **Tip:** Use the **Find** field to locate a host in a long list.

5. To remove a host, select the text and press delete.

6. Click **Save** to save your changes.

Additional Information

- The syntax of your host entries are not validated until you try to save the host list. If you entered an invalid host specification, the Host List Editor indicates which host is invalid so that you can correct the problem and try saving again. If there is an issue with multiple hosts, the Editor alerts you to them one at a time.

- The Host List Editor does not validate the existence of the hosts.

Delete a Host List in the Sky UI

You can use the Host List Editor to delete a host list that is no longer needed. Note that the Editor prevents you from deleting a host list that is currently specified in a matching rule.

To delete a host list:

1. Make sure there aren't any classes using the host list you want to delete.

2. Open the Host List Editor. **show screen...**

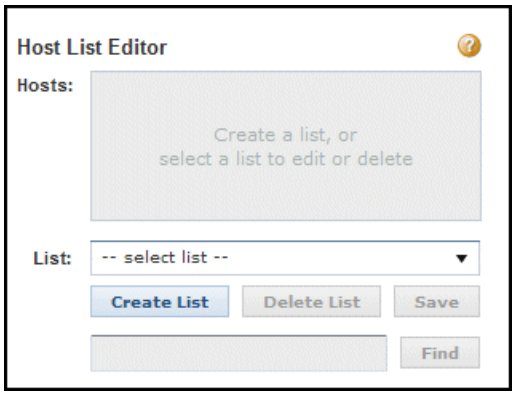

3. In the **List** field, select the host list to delete. The hosts associated with this list are displayed in the **Hosts** text box.

4. Click **Delete List** and click OK to confirm.

> **Note:** If you try to delete a host list that is currently specified in a matching rule, you will see the following message: *Hostlist xxxx is in use*. Before you can delete this host list, you need to determine which classes specify this host list in a matching rule.

Sky UI Class Tree Management Tasks

To help you maintain and organize your traffic tree, the Sky UI offers tools for grouping classes into folders, moving classes to different branches in the tree, deleting and copying classes, and using class lists.

Configure Class Tree Settings

Several settings are available to configure class tree settings in the Sky UI.

To configure class tree settings:

1. Click **Traffic Management**. The class tree appears in the upper pane. **show screen...**

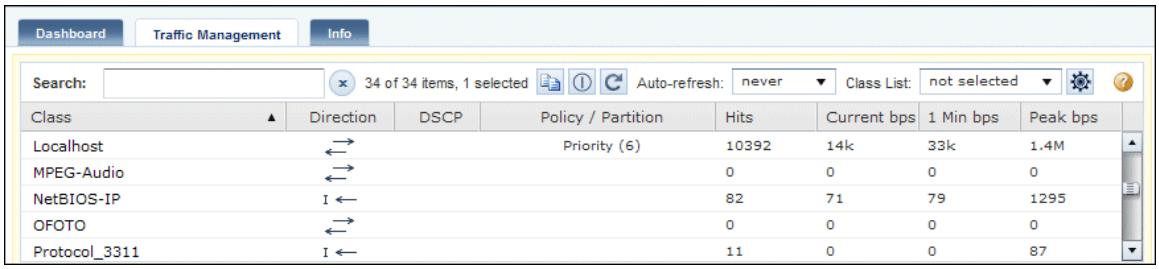

2. Click the **Settings** icon and make sure the **Class Tree** tab is selected. The settings appear in the lower pane, below the class tree. **show screen...**

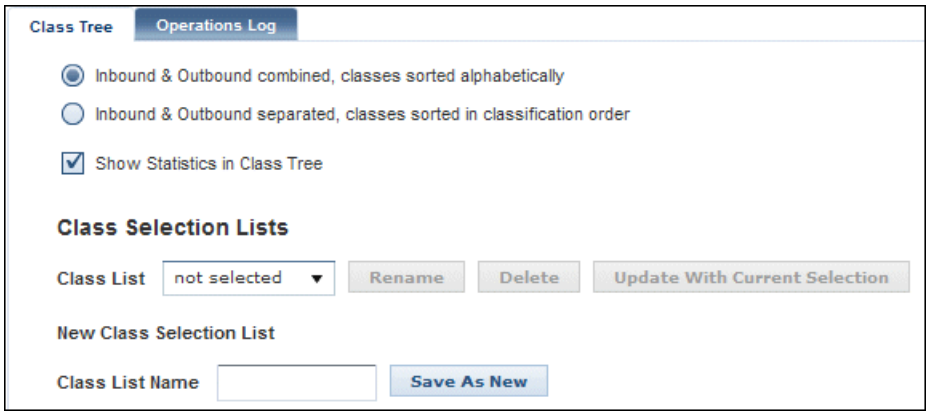

3. Enable/disable the following settings:

Option	Description
Inbound & Outbound combined	When a class exists in both the Inbound and Outbound direction, the two classes are combined into a single class in the tree, and a Direction column appears to indicate whether the class is for Inbound (I ←), Outbound (□ →), or both (⇄). The classes are listed in alphabetical order. **show screen...**

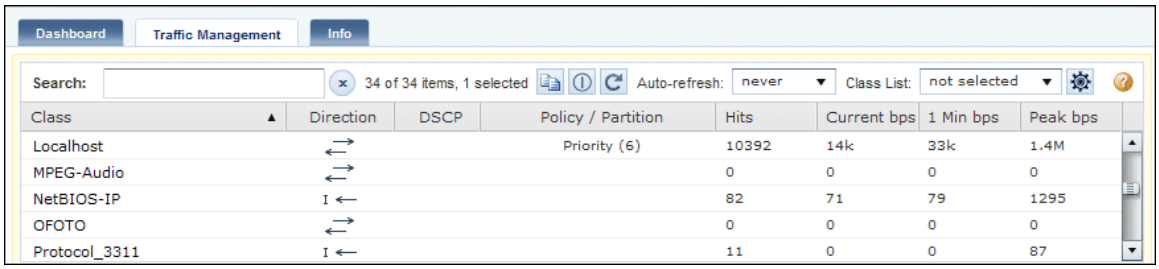

Note that statistics are also combined (for example, Hits represents the total class hits for Inbound and Outbound). Combining Inbound and Outbound classes presents a streamlined tree and offers an easier approach to class/policy management. For example, if you want to apply the same policy to Inbound and Outbound, you only have to select one class.

Note: Inbound and Outbound classes with the same name, policy, and partition are shown as a combined class even if they have different matching rules. However, if the policy or partition of the Inbound and Outbound classes are not identical, they are shown as two separate classes in combined view.

Inbound & Outbound separated	Inbound and Outbound classes appear separately in the class tree. All the Inbound classes are presented first, followed by the Outbound classes. The class tree is presented this way in the Advanced UI. The Direction column is unnecessary. **show screen...**

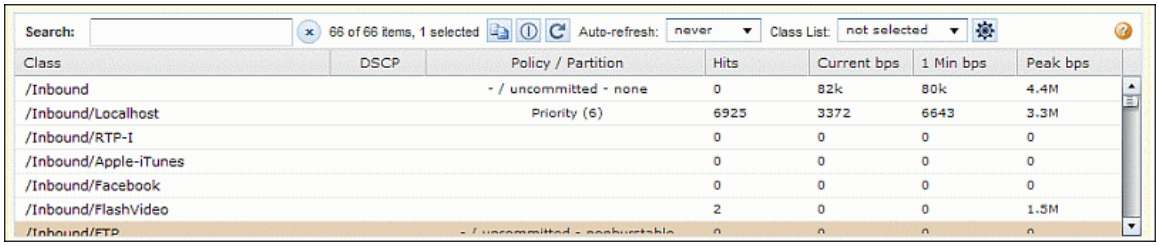

With this option, the class tree is sorted in classification order: classes are listed in order from most specific to least specific. Classes with more specific criteria are at the top of the tree. PacketShaper classifies a traffic flow by searching the class tree, beginning at the top, to find a match for the flow. As soon as PacketShaper finds a match, it categorizes the flow into the appropriate class and looks no further down the tree (even if there may be another class that matches the traffic flow).

Option	Description
Show Statistics in Class Tree	**Enabled:** Displays the columns of statistical data: Hits, Current bps, 1 Min bps, Peak bps. **show screen...**

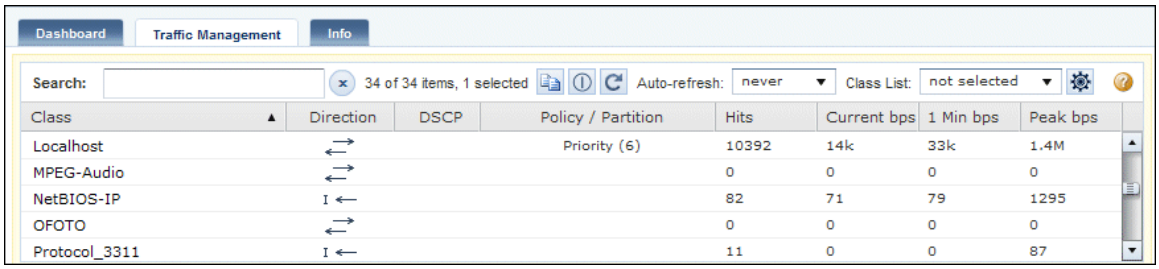

Disabled: Hides the columns of statistical data. **show screen...**

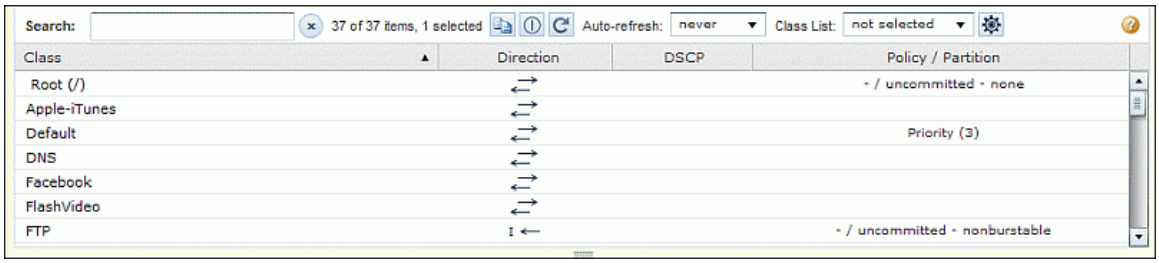

The **Show Statistics** setting is active for the current browser session only.

Add a Folder in the Sky UI

You can use a folder to organize the class tree. Two folders are created automatically: /Inbound and /Outbound. In addition, if port-based traffic is autodiscovered, the classes for this traffic are placed in DiscoveredPorts folders at the bottom of the /Inbound and /Outbound sections of the class tree.

Folders do not match any traffic that is, there are no specific traffic criteria associated with this type of class. Therefore, a policy cannot be applied to a folder. You can, however, apply a partition to a folder.

To add a folder:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

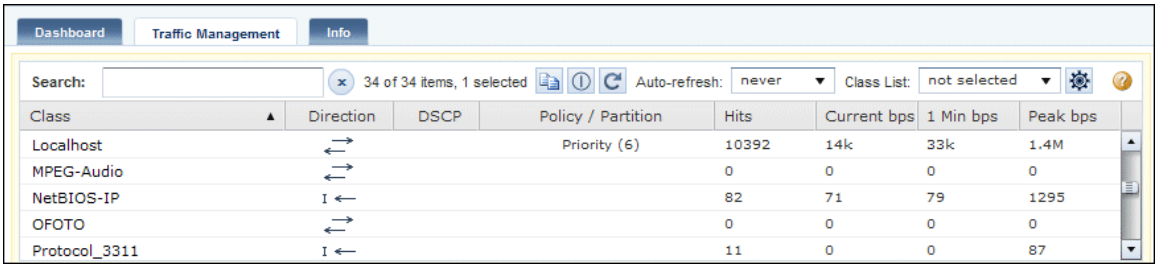

2. Select the parent class for the folder.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Add Folder**. **show screen...**

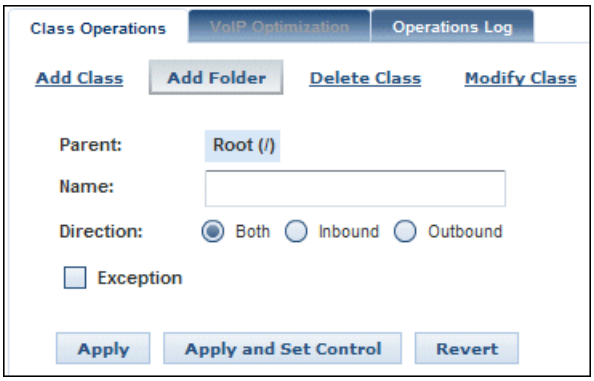

6. In the **Name** field, type the folder name (up to 31 characters). Use only alphanumeric characters and underscores (_), dashes (-), and periods (.).

7. Select the **Direction** for the new class: **Both**, **Inbound**, or **Outbound**

8. To create an exception class, enable the **Exception** checkbox.

9. Click **Apply**.
or
If you want to create the folder and immediately assign a partition to it, click **Apply and Set Control**.

> **Note:** A folder icon appears to the left of the folder name in the class tree.

Delete Classes in the Sky UI

When you delete a class, the class name is removed from the class tree and PacketShaper will no longer classify flows for the associated service or application. Keep in mind, though, that if traffic discovery is enabled, the discovery process may recreate the class based on the traffic it discovers.

To delete classes from the class tree:

1. Click **Traffic Management**. The class tree appears in the upper pane. **show screen...**

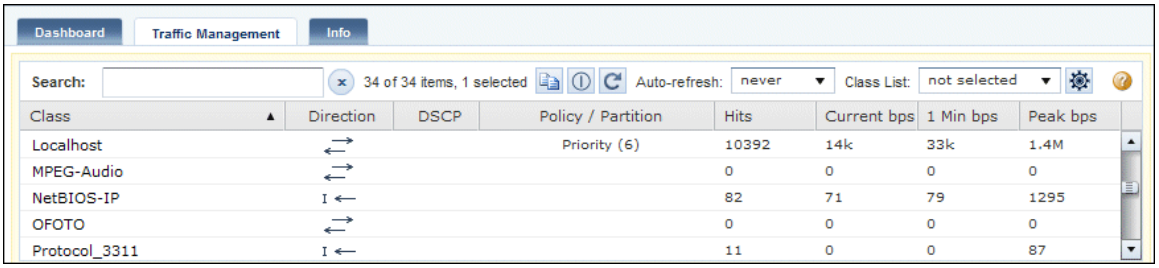

2. Select the classes you want to delete (up to 10), using any of the following techniques:
   - Click anywhere on the class row to select a single class.
   - Shift+click to select a range of classes.
   - Ctrl+click to select non-adjacent classes.
   - Use the **Search** field to locate the class(es) and then select each one.
3. Click the **Policy Manager** icon. Policy Manager appears in the lower pane, below the class tree.
4. Click the **Class Operations** tab.
5. Click **Delete Class**. If you had selected a single class, its name appears. If you selected multiple classes, the screen indicates the number of classes that will be deleted (for example, *Delete 4 Selected Classes*). **show screen...**

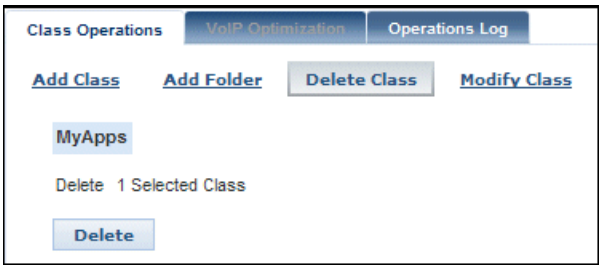

6. Click **OK** to confirm the deletion.

Move Classes in the Sky UI

You can move classes into a folder or to other branches in the traffic tree.

To move one or more classes:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

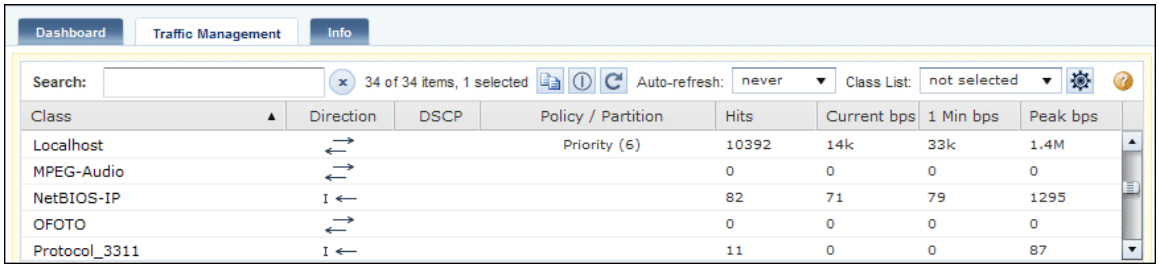

2. Select the class(es) you want to move.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Modify Class**. **show screen...**

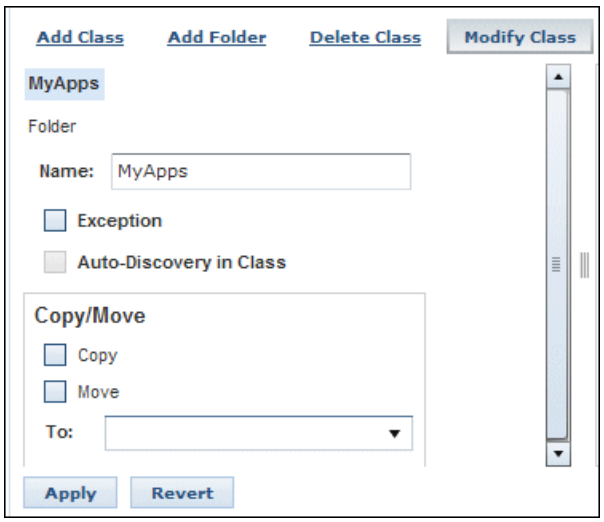

6. Select the **Move** checkbox.

7. In the **To** list, select the parent class or folder where you want to move the class(es).

8. Click **Apply**.

Copy Classes in the Sky UI

You can copy classes into a folder or to other branches in the traffic tree.

To copy one or more classes:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

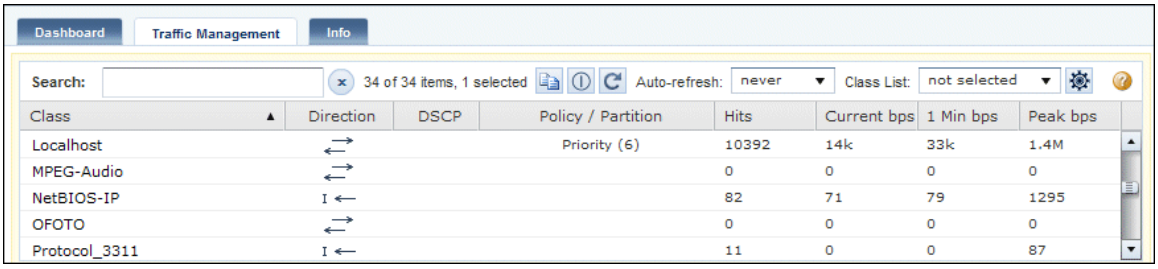

2. Select the class(es) you want to copy.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Modify Class**. **show screen...**

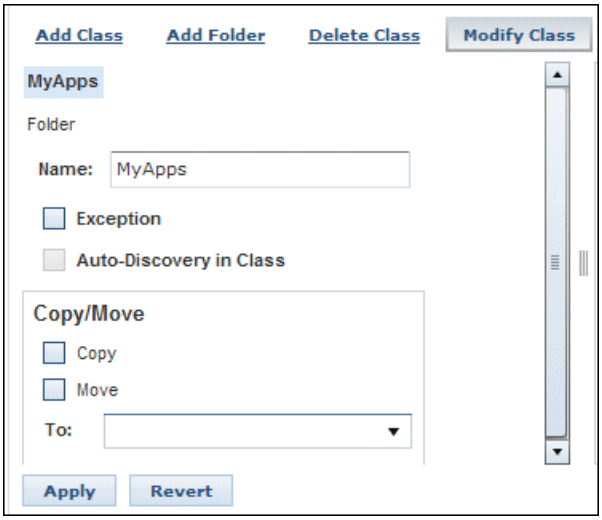

6. Select the **Copy** checkbox.

7. In the **To** list, select the parent class or folder to where you want to copy the class(es).

8. Click **Apply**.

Rename a Class in the Sky UI

You are allowed to rename any class in the tree, as long as the new name is unique and follows naming conventions. Note that when renaming a class you cannot change just the case; for example, you cannot rename HTTP to http. If the selected class is a combined class, Sky will automatically rename the class in both directions (Inbound and Outbound).

To rename a class:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

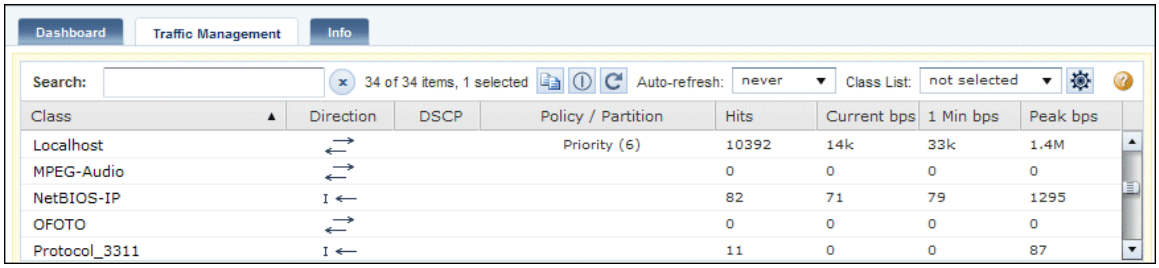

2. Select the class you want to rename.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Modify Class**. **show screen...**

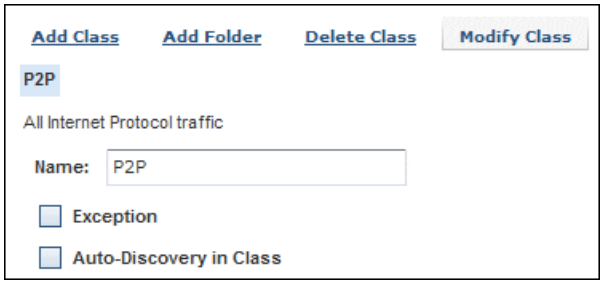

6. In the **Name** field, type the new class name (up to 31 characters). Use only alphanumeric characters and underscores (_), dashes (-), and periods (.).

7. Click **Apply**.

Create Class Lists in the Sky UI

The Sky UI allows you to assign a name to a set of classes. Once the list name is created, you can quickly select the set of classes by choosing the name. If you find yourself frequently selecting the same group of classes for class/policy management purposes, the class list feature allows you to speed up the selection.

To create a class list:

1. Click **Traffic Management**. The class tree appears in the upper pane. **show screen...**

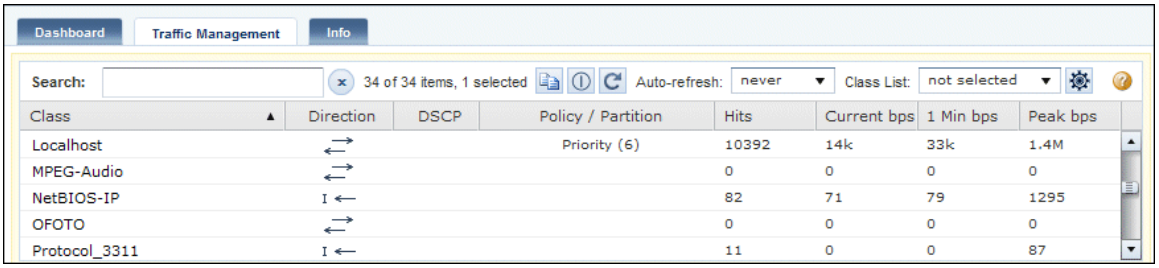

2. Select all the classes you want in your list, using any of the following techniques:

   - Click anywhere on the class row to select a single class.
   - Shift+click to select a range of classes.
   - Ctrl+click to select non-adjacent classes.
   - Use the **Search** field to locate the class(es) and then select each one.

3. Click the **Settings** icon and make sure the **Class Tree** tab is selected . The settings appear in the lower pane, below the class tree. **show screen...**

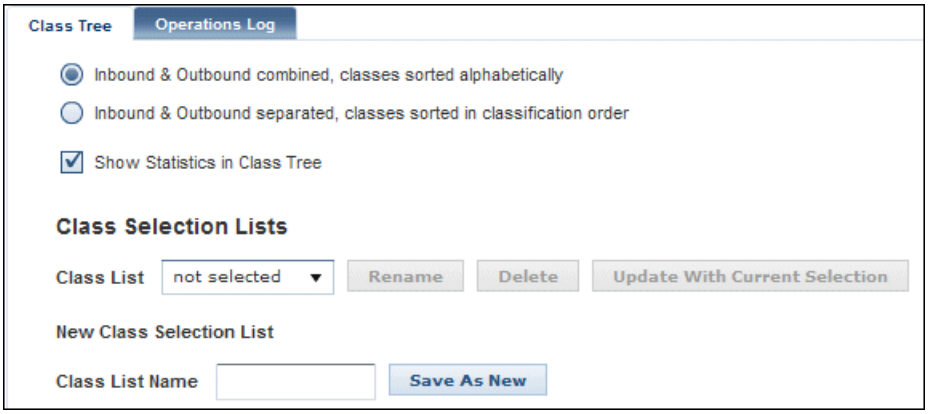

4. In the **Class List Name** field, enter the name for your class list. The name can be up to 31 characters long; alphanumeric values are allowed, including spaces, hyphens, and underscores.

5. Click **Save as New**.

Additional Information

- A class list can contain up to 10 classes.
- You can create up to 50 class lists.

Use Class Lists in the Sky UI

Once a class list name is created, you can quickly select the set of classes by choosing the name from the **Class List** drop-down field in the class tree pane.

To select a class list:

1. Click **Traffic Management**. The class tree appears in the upper pane. **show screen...**

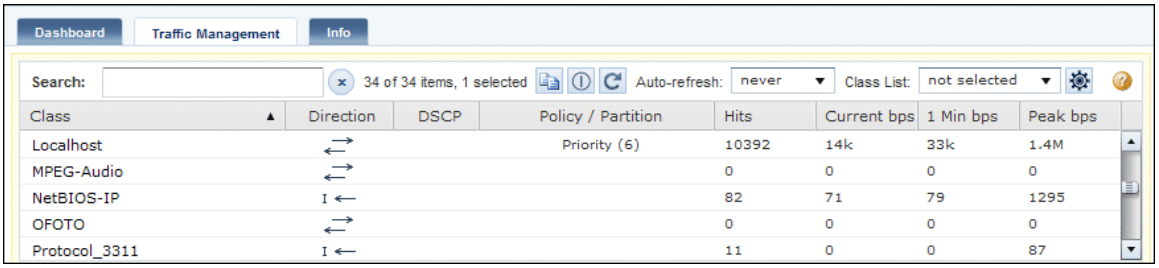

2. Select the class list name from the **Class List** drop-down field. The classes associated with the list are selected in the class tree; you may need to scroll to see the classes–they aren't necessarily grouped together.

3. Perform the desired class/policy operation on the selected classes.

> **Note:** After selecting a class list, you can confirm the number of selected classes by looking at the top of the class tree pane (for example: *939 of 939 items, **6 selected***).

Modify a Class List in the Sky UI

After creating a class list, you may want to add/remove classes to/from the selection.

To modify a class list:

1. Click **Traffic Management**. The class tree appears in the upper pane. **show screen...**

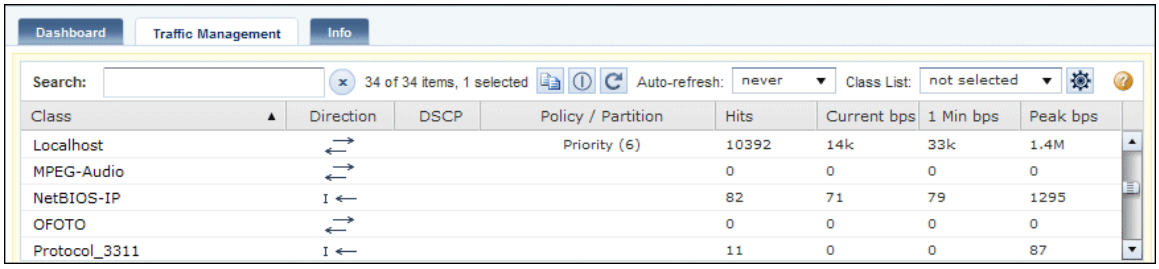

2. Select the class list name from the **Class List** drop-down field. The classes associated with the list are selected in the class tree.

3. To add another class to the class list, hold down Ctrl as you click the class name.

4. To remove a class from the class list, hold down Ctrl as you click the class name. The class row will no longer be selected.

5. Click the **Settings** icon and make sure the **Class Tree** tab is selected. The settings appear in the lower pane, below the class tree. **show screen...**

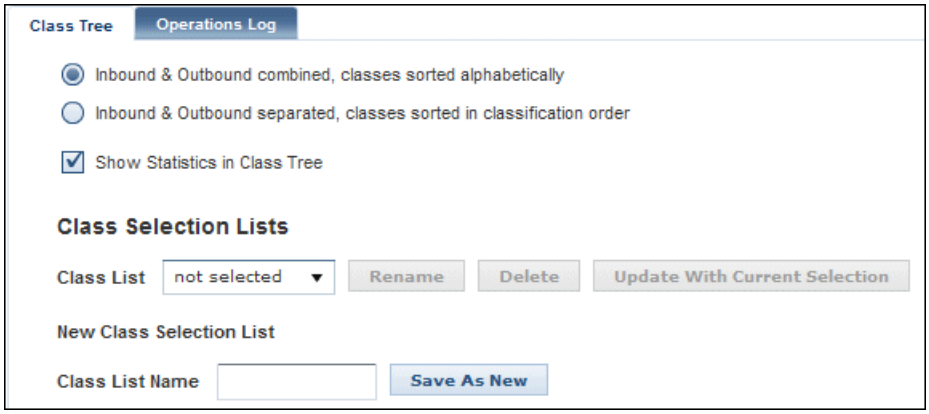

6. In the **Class List** drop-down list, select the class list name.

7. Click **Update with Current Selection**. **show screen...**

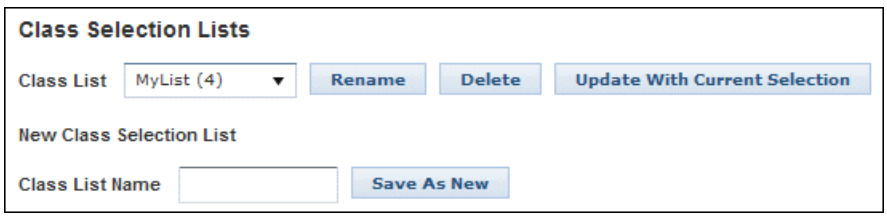

Additional Information

- After selecting a class list, you can confirm the number of selected classes by looking at the top of the class tree pane (for example: *939 of 939 items, **6 selected***).
- To delete a list, select its name from the **Class List** and click **Delete**. Note that this deletes the class list name–not the traffic classes associated with the name.

Rename a Class List in the Sky UI

To rename a class list:

1. Click **Traffic Management**.
2. Click the **Settings** icon. The settings appear in the lower pane, below the class tree.
3. Click the **Class Tree** tab.
4. In the **Class List** field, select the class list name. **show screen...**

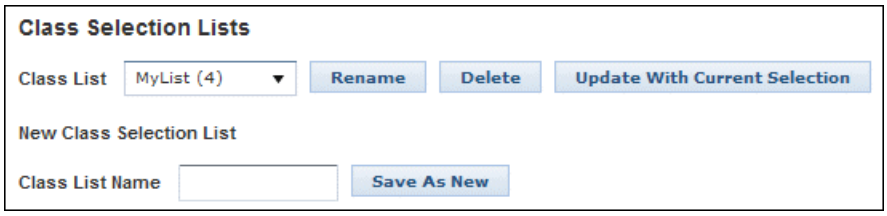

5. Click **Rename**. A new, empty field appears for you to enter the new name.
6. Enter the new class list name (up to 31 characters).
7. Click **Save**.

Copy the Class Tree to the Clipboard

To save a picture of the class tree along with its statistics, you can copy the class tree to the Windows clipboard, paste it into Microsoft Excel, and save it in a file.

To copy the class tree to the clipboard:

1. Click **Traffic Management**. The class tree appears in the upper pane. **show screen...**

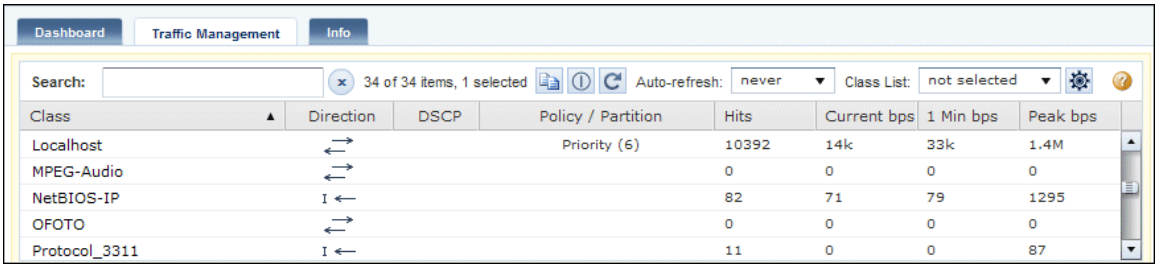

2. Make sure the class tree displays the classes you want to copy. If necessary, clear the **Search** field.

3. Click the icon above the class tree. The classes, settings, and statistics are copied to the clipboard.

4. Open Microsoft Excel and paste the clipboard contents into an empty spreadsheet.

The class tree data is copied into separate spreadsheet cells. Adjust column widths as necessary, and save the file.

Select Which Columns and Rows to Display in the Sky UI Class Tree

The Sky UI class tree, by default, lists all classes in your traffic tree, and provides a variety of statistics (such as class hits, current rate, and peak rate) and settings (such as policies and partitions) associated with each of these classes. However, you have the ability to select which classes are displayed as well as selectively list and sort columns. This capability gives you the power to zoom in on the information you need.

Hiding Statistics

If you don't want to see the real-time statistics in the class tree, you can hide them. To suppress the display of the statistics:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

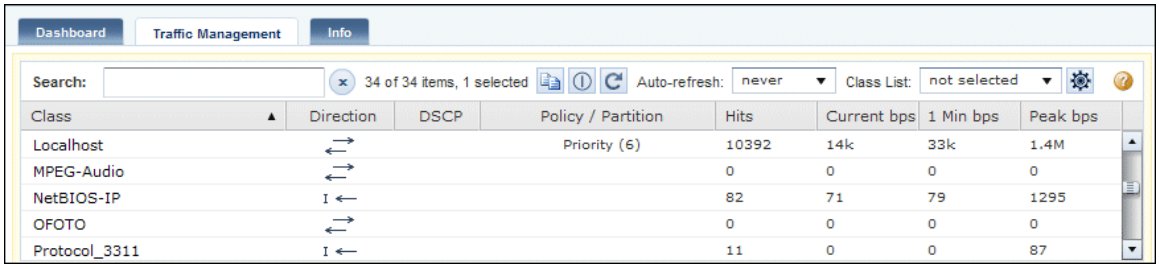

2. Click the Settings icon – either the one above the class tree or the one on the lower left part of the screen. **show screen...**

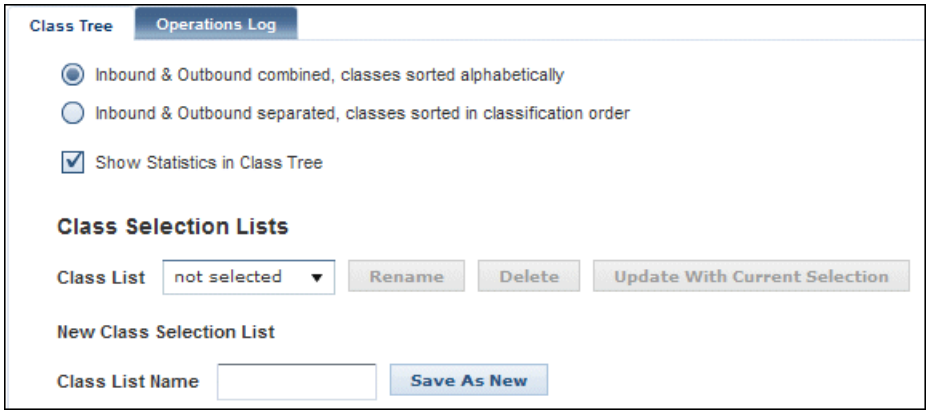

3. Unselect the **Show Statistics in Class Tree** checkbox. The Hits, Current bps, 1 Min bps, and Peak bps columns are now hidden. **show screen...**

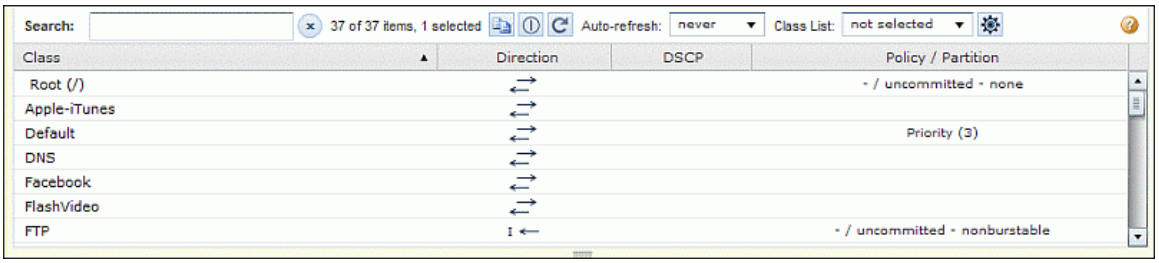

Note that this setting is temporary and is good for the current session only. The next time you log in, the statistics will display in the class tree.

Sorting the Class Tree

The order of the classes in the class tree depends on whether the Inbound and Outbound classes are combined or separated. If they are combined, the classes are listed in alphabetical order. If they are separated, the classes are sorted in classification order, from most to least specific; classes with more specific criteria are placed higher on the list. For example, a host-based class for a specific IP address would appear above the more general service-based (application) classes. Classes that are the same relative level of specificity are alphabetized. See Configure Class Tree Settings for steps on specifying your class tree preference (Inbound/Outbound combined versus separate).

You can change the sort order by clicking a column heading in the class table; you can sort in ascending or descending order. For example, if you click the Hits heading, the class tree is rearranged according to the number of hits a traffic class has incurred, with the least hits first. If you click the Hits heading again, the tree is sorted by class hits in descending order (from most to least).

An icon next to the heading indicates the sort column:

- ▲ indicates ascending order
- ▼ indicates descending order

Column widths can be adjusted by dragging the separator line between column headings.

Filtering the Class Tree

With auto-discovery enabled, your class tree can contain hundreds or thousands of applications. To narrow down this list, you can filter the class tree by entering a search string in the **Search** field. Only classes containing the string will be displayed. For example, you can enter *http* to see all classes with http in the name. The **Search** field also accepts the | symbol which indicates an "or" operation (for example, *http|ftp* to display all http or ftp classes).

When the class tree is filtered, only the classes that match the search criteria are displayed. To the right of the Search field, it indicates how many classes were found (such as *2 of 108 items*).

To clear the **Search** field, click the X icon next to field. All traffic classes will now be displayed in the class tree.

Update the Sky UI Class Tree Statistics

The statistics in the Sky UI class tree are captured the moment you click the **Traffic Management** tab. To get a new snapshot, click the **Refresh Class Tree Now** icon. This button updates the statistics for existing classes and displays any newly discovered traffic classes.

You can set the class tree to automatically update at regular intervals. To enable auto-refresh:

1. Display the **Auto-refresh** drop-down list.

2. Select one of the listed intervals. The choices are between 1 minute and 1 hour.

To stop automatic updating, select **never** from the **Auto-refresh** list.

Reset the Sky UI Class Tree Statistics

The following real-time statistics are listed in the Sky UI class tree: **show screen...**

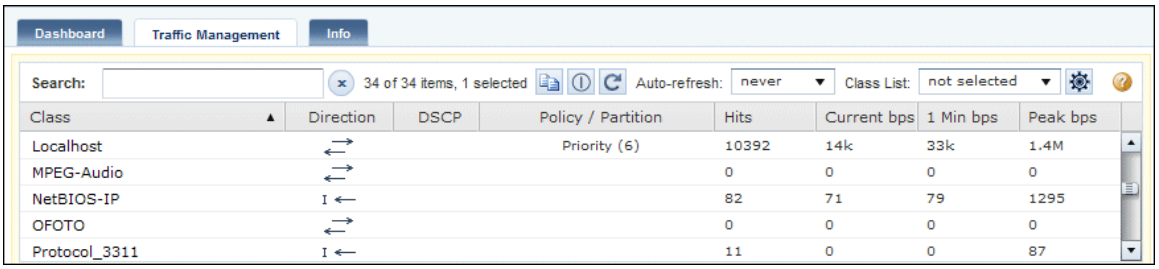

- Hits
- Current bps
- 1 Min bps
- Peak bps

To reset the real-time statistics, click the **Reset Values** icon. All values on the screen are zeroed out. Occasionally, some statistics do not have zero values, even after you click the icon; this happens because PacketShaper may be logging traffic flows during the clearing of statistics.

Sky UI Traffic Control Tasks

PacketShaper offers several techniques for controlling traffic on your network: policies, partitions, and DSCP markings. Note that traffic shaping must be enabled in order for these controls to take effect.

Control Your Traffic

PacketShaper offers three techniques for controlling your traffic to ensure quality of service (QoS):

- **Policies** A policy determines how a class's *individual flows* are treated in the context of competing traffic, and allows you to manage bandwidth on a flow-by-flow basis. With policies, you can give each flow of mission-critical traffic the bandwidth it needs for optimum performance, as well as protect it from greedy, less important traffic. In addition, policies can keep non-urgent traffic flows (such as FTP) from consuming more than an appropriate share of bandwidth.
- **DSCP** Network routers can apply differentiated grades of service to packet streams based on a Diffserv Code Point (DSCP) marker in an IP packet header. PacketShaper is able to mark applications with a specific DSCP marker (0-63).
- **Partitions** A partition manages bandwidth for a traffic class' *aggregate flows*, so that all of the flows for the class are controlled together as one. You can use partitions to protect mission-critical traffic by guaranteeing that a traffic class always gets a defined amount of bandwidth. In addition, you can use partitions to limit aggressive, non-critical traffic by allowing that traffic class to consume only a defined amount of bandwidth.

> **Note:** Policies, DSCP, and partitions do not take effect unless shaping is enabled.

To set controls on one or more traffic classes:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane.**show screen...**

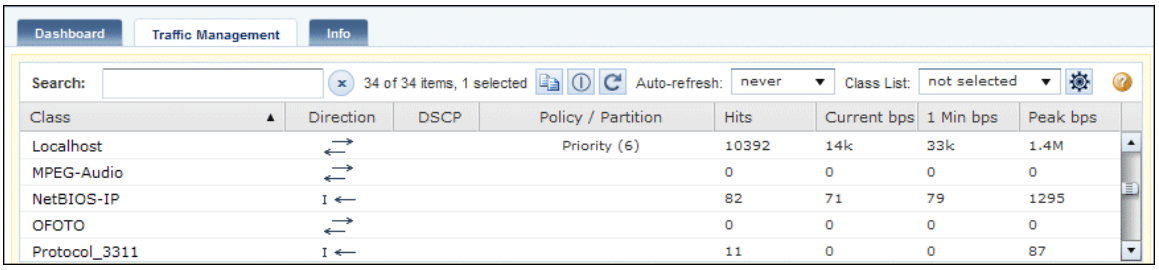

2. Select the class(es) for which you want to set the same controls.
3. Click the **Policy Manager** icon.
4. Click the **Class Operations** tab.

5. Click **Control Traffic**. **show screen...**

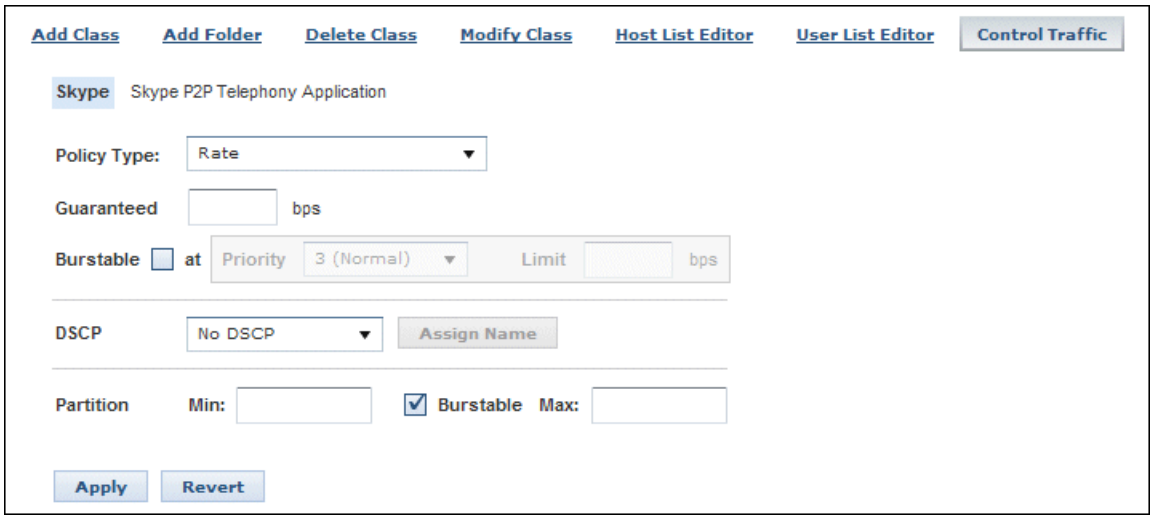

6. Apply one or more of the following control techniques:

   - Assign a policy. See Policy Overview.
   - Select a DSCP value. See DSCP Marking.
   - Create a partition. See Partition Overview.

7. Click **Apply**.

Create Partitions in the Sky UI

To define/modify a partition for one or more traffic classes:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

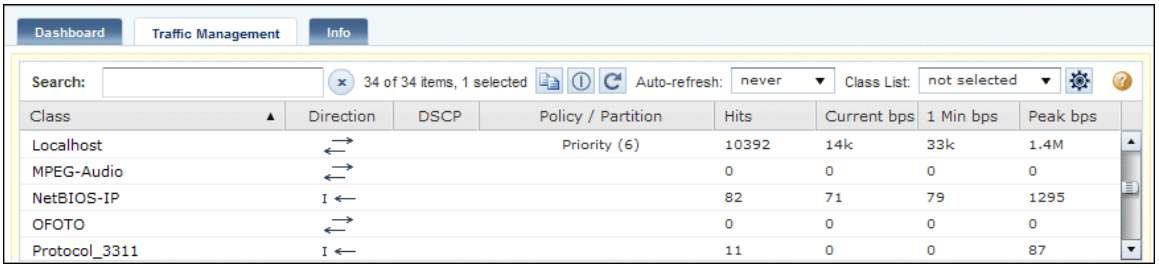

2. Select the class(es) for which you want to assign the same partition size.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Control Traffic**. **show screen...**

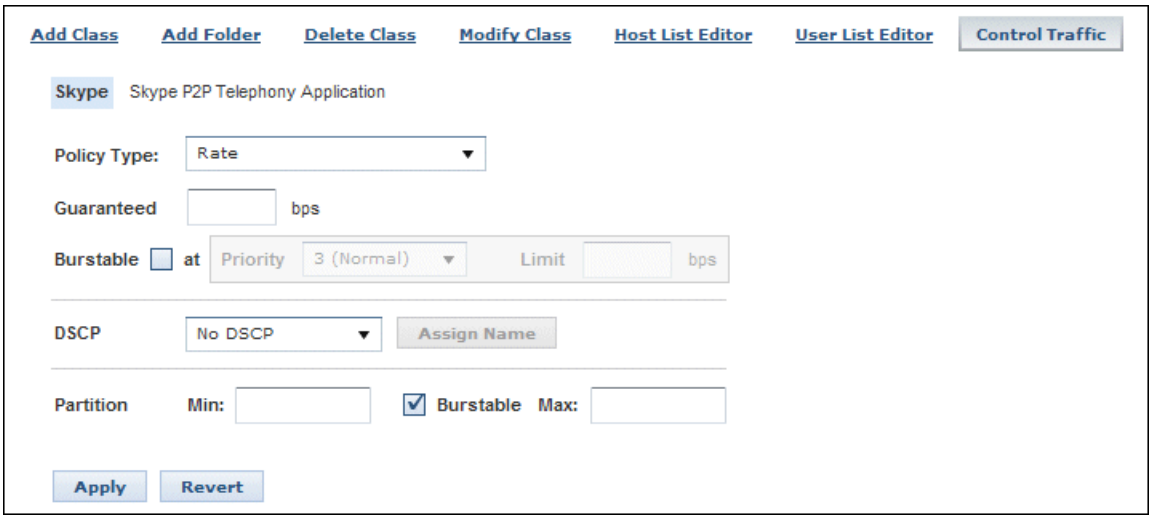

6. For the partition **Min**, enter the minimum partition size, expressed in bits per second (bps) or as a percentage of the parent partition's size (%). If the size is entered in bits per second, you can follow the integer by a “k” (thousands), “M” (millions), or "G" (billions). The minimum partition size is 1024 bps. The maximum size varies by model.

PS-S200: licensed bandwidth

PS-S400: 2 Gbps or licensed bandwidth (whichever is less)

PS-S500: 2 Gbps. Exception: The Inbound and Outbound partitions can have a maximum size of 10 Gbps if you have a 10G license key or 5 Gbps if you have a 5G key.

> **Note:** Not all classes have partitions, so when you use the % option, it will be a percent of the nearest parent class with a partition (which may be /Inbound or /Outbound or may be some intermediate class if your tree has multiple levels of hierarchy).

7. (Optional) Select the **Burstable** checkbox to allow the partition to use available excess bandwidth if needed. When this box is unchecked, the partition has a fixed size; when it is not using its reserved bandwidth, that bandwidth is available to other traffic.

8. (Optional) If you selected **Burstable** in step 7, you can enter a value in the **Max** field to cap the amount of excess bandwidth for the partition. If you don't enter a value, the partition size has no maximum limit. The value can be specified in bits per second or as a percentage of the parent partition's size.

9. Click **Apply**.

> **Note:** If you get an error message, you can look at the operation log.

Set a Rate Policy in the Sky UI

Using a *rate policy*, you can establish a minimum rate for each flow of an application, allow the flow prioritized access to excess bandwidth if it's available, and set a limit on the total bandwidth that a flow can consume. A rate policy is especially suited for traffic that tends to burst.

Rate policies are recommended for applications such as web traffic (HTTP), secure web traffic (SSL), email (SMTP and POP3), file transfer (FTP), newsgroups (NNTP), news (EntryPoint), and streaming media (RealAudio, MPEG-Audio, and MPEG-Video). For more information, see Rate Policy Examples.

For latency-sensitive applications such as Voice over IP (VoIP), you can configure a policy to allocate sufficient bandwidth to assure quality of service. That is, no matter what else is competing for access-link bandwidth, VoIP will get the minimum amount of bandwidth it needs, assuming that there aren't too many active VoIP sessions already. For details, see VoIP Optimization.

A rate policy enables precise management of TCP traffic. For User Datagram Protocol (UDP) traffic, a rate policy uses delay-bound scheduling. Rate policies are *not* recommended for non-IP traffic.

Setting a Rate Policy

The most common rate policy definition uses the following settings:

- Guaranteed Rate = 0 bps
- Burstable at a priority (for a ranking relative to other applications)
- Limit the amount of excess bandwidth (if the burstable setting is used)

To apply a rate policy to one or more traffic classes:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

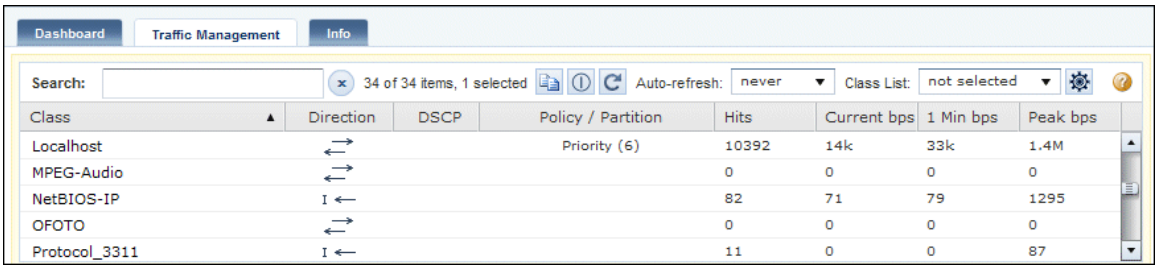

2. Select the class(es) for which you want to assign a priority policy.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Control Traffic**. **show screen...**

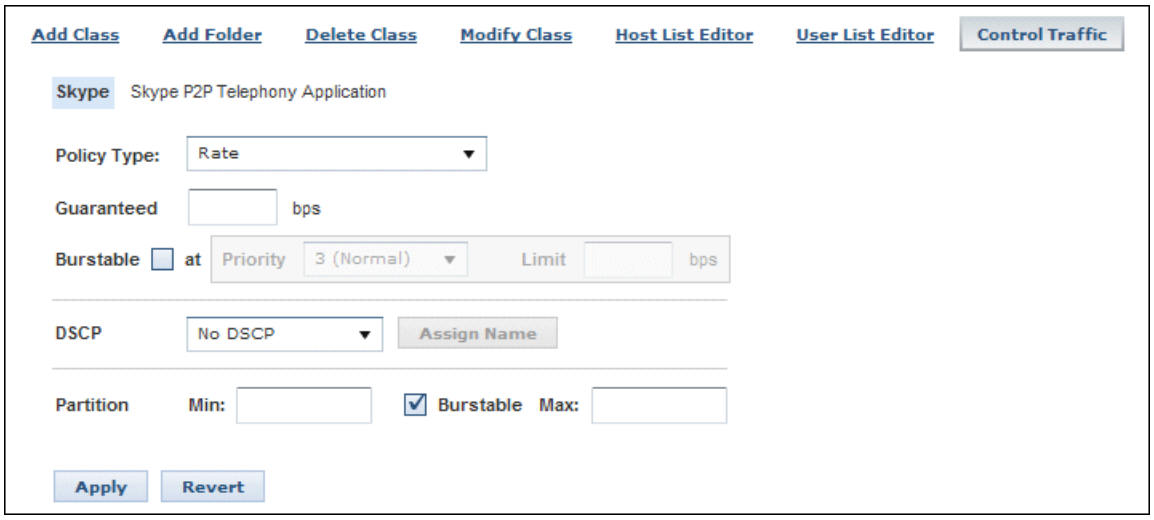

6. From the **Policy Type** drop-down menu, select **Rate**.

7. In the **Guaranteed** field, enter the minimum guaranteed rate given to each connection. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or "G" (billions). When a specific minimum rate is not required, set the rate to 0 bps and configure the **Burstable** option. The guaranteed rate must be a minimum of 1024 bps.

8. If you want the class to be able to access unused bandwidth (that is, to be *burstable*), select the **Burstable** checkbox and select a **Priority** (0 through 7) from the pull-down menu. Note that 0 is the lowest priority and 7 is the highest. The priority level defines how the class will compete with other classes for excess bandwidth.

9. (Optional) To cap the excess rate available to the selected class(es), enter a value in the **Limit** field. If a limit is specified, it must be a minimum of 256 bps and must be greater than the policy's guaranteed rate.

> **Note:** Remember that a policy controls how much bandwidth is allocated to each traffic flow. Do *not* assign large amounts of guaranteed bandwidth to a traffic class because, as the number of simultaneous sessions increases, you'll run out of bandwidth.

10. Click **Apply**.

> **Note:** If you get an error message, you can look at the operation log.

To see how bandwidth is being managed for the class(es), you can create reports. Once you have a good understanding of the number of sessions on the network and any requirements to guarantee a specific service level, you can start adjusting the amount of guaranteed bandwidth to get the desired result.

Set a Priority Policy in the Sky UI

Priority policies are used to establish a priority for traffic without specifying a particular rate. Use a priority policy for short, non-continuous traffic flows or traffic that does not burst. The priority determines the importance of the traffic class; you specify a value from 0 to 7, where 7 is the highest priority. The Inbound/Default and Outbound/Default classes are factory-set to a priority of 3.

For example, you might define a priority 7 policy for TN3270 traffic to ensure quality of service for critical interactive sessions. Then, to keep print traffic from taking over the link, assign a priority 2 policy to the TN3287 traffic class. Finally, give the Quake game a priority of 0.

Setting a Priority Policy

To apply a priority policy to one or more traffic classes:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

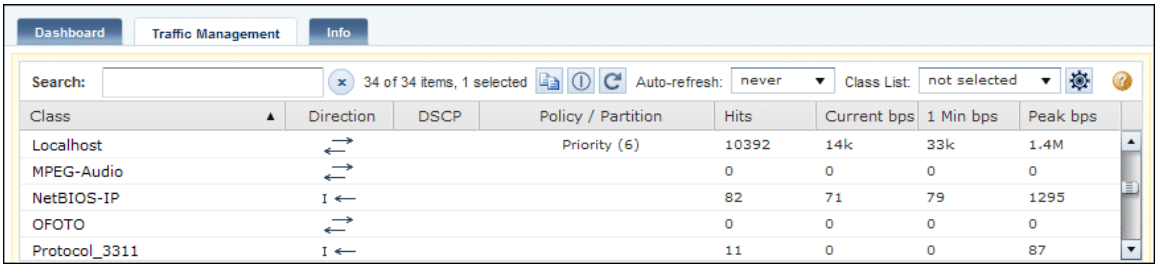

2. Select the class(es) for which you want to assign a priority policy.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Control Traffic**. **show screen...**

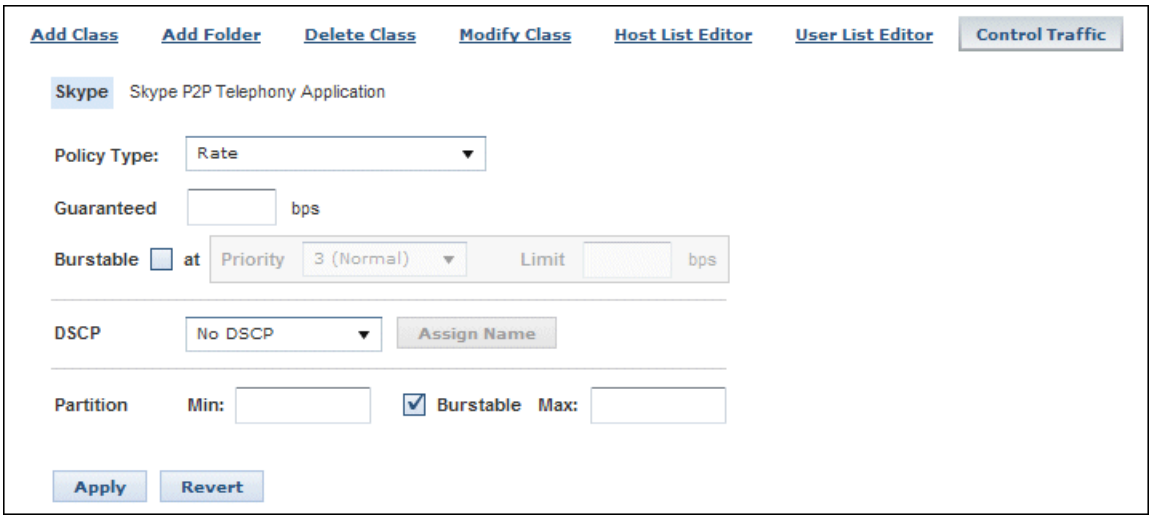

6. From the **Policy Type** drop-down menu, select **Priority**.

7. Select a **Priority** (0 through 7) from the pull-down menu. Note that 0 is the lowest priority and 7 is the highest.

8. Click **Apply**.

> **Note:** If you get an error message, you can look at the operation log.

Priority Policy Examples

Use priority policies to prioritize specific applications that need preferential or low-priority treatment. The following table lists some typical uses of priority policies.

Traffic Class	Recommended Priority Level	Explanation
DNS LDAP	6	Directory lookups typically are very small, so let them through quickly to start the transaction.
SNMP	2	Network management flows usually do not burst and are not time sensitive.
TN3270 TN5250	7	These host-access applications have small, interactive, time-sensitive flows.

Set a Never-Admit Policy in the Sky UI

A *never-admit policy* tells PacketShaper to enforce admission control at the beginning of each flow; the connection willl be refused. Note that this policy type is for TCP traffic only; for UDP traffic, use a discard policy.

> **Note:** When shaping is turned off, a class with a never-admit policy will display a rate of zero in the class tree (Traffic Management tab). Packets are passed through PacketShaper, but the rate values make it appear that the traffic is not passing through the unit.

To apply a never-admit policy to one or more traffic classes:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

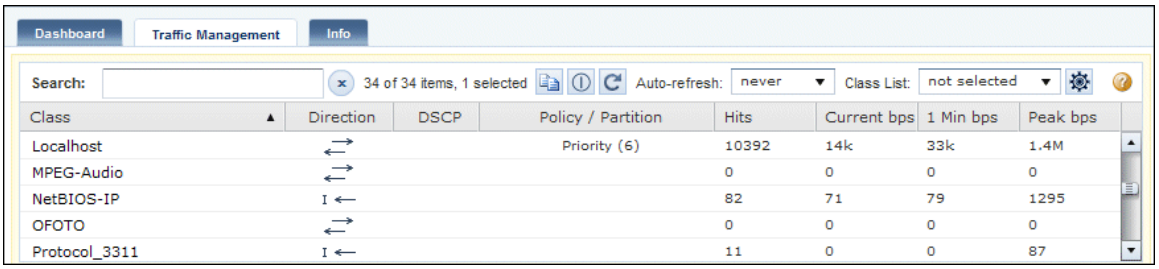

2. Select the class(es) for which you want to assign a never-admit policy.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Control Traffic**. **show screen...**

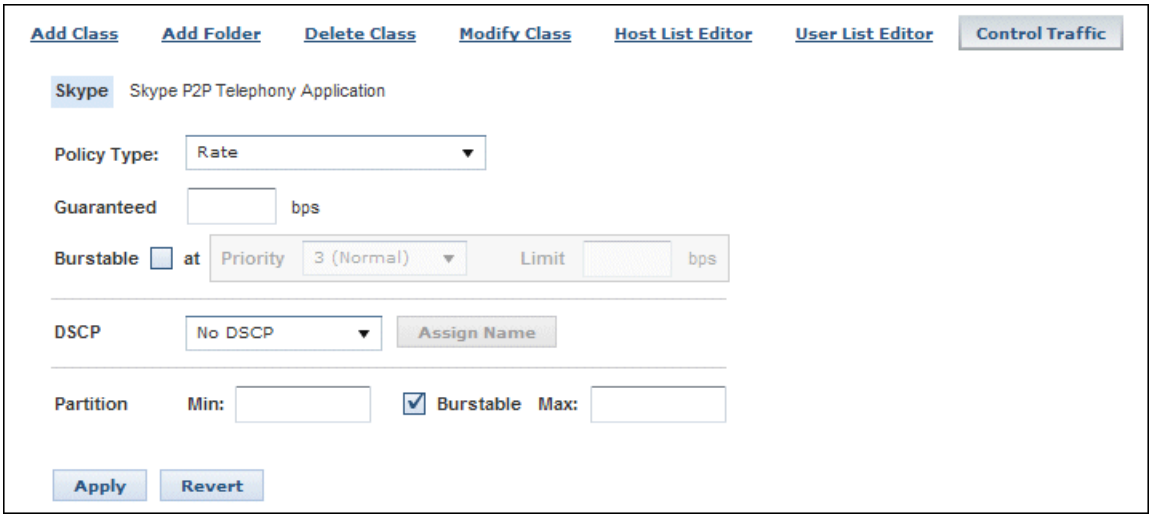

6. From the **Policy Type** drop-down menu, select **Never-Admit**.

7. For web traffic, select one of the **Admission Control** options below:

Web Refuse to refuse the connection
-or-
Web Redirect to redirect the user to another web page

8. If you selected Web Redirect, specify the URL in the **Redirect URL** text box. Specify the full URL – for example, *http://www.mycompany.com/main.htm*.

9. Click **Apply**.

> **Note:** If you get an error message, you can look at the operation log.

Set a Discard Policy in the Sky UI

Use a *discard policy* to toss all packets for a traffic class. You may want to block a service or application because it is:

- Consuming too much of your network bandwidth
- Considered non-essential to your business
- Prohibited on your network

A discard policy is recommended for blocking UDP traffic classes because UDP traffic is not session oriented. For TCP or web traffic, you would typically use a never-admit policy to block traffic and inform the user that the session was disallowed. If you use a discard policy for TCP traffic, the discarded packets will cause TCP timeouts resulting in a long wait period before the user gets any feedback. In some cases this may be desirable. For example, you may want to "time out" someone attempting to crack your web server.

To apply a discard policy to a traffic class:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

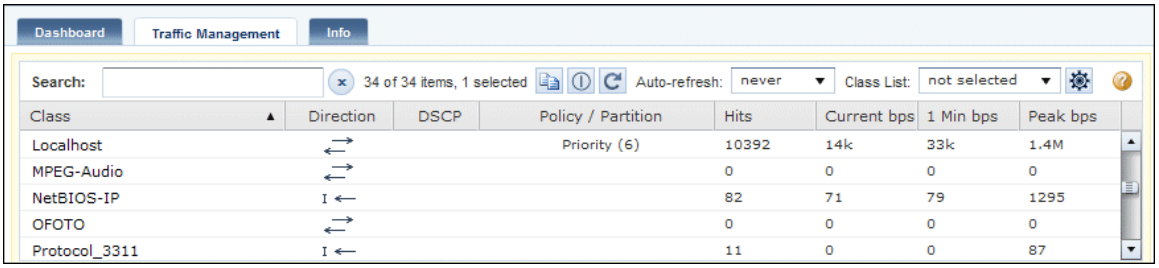

2. Select the class(es) for which you want to assign a discard policy.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Control Traffic**. **show screen...**

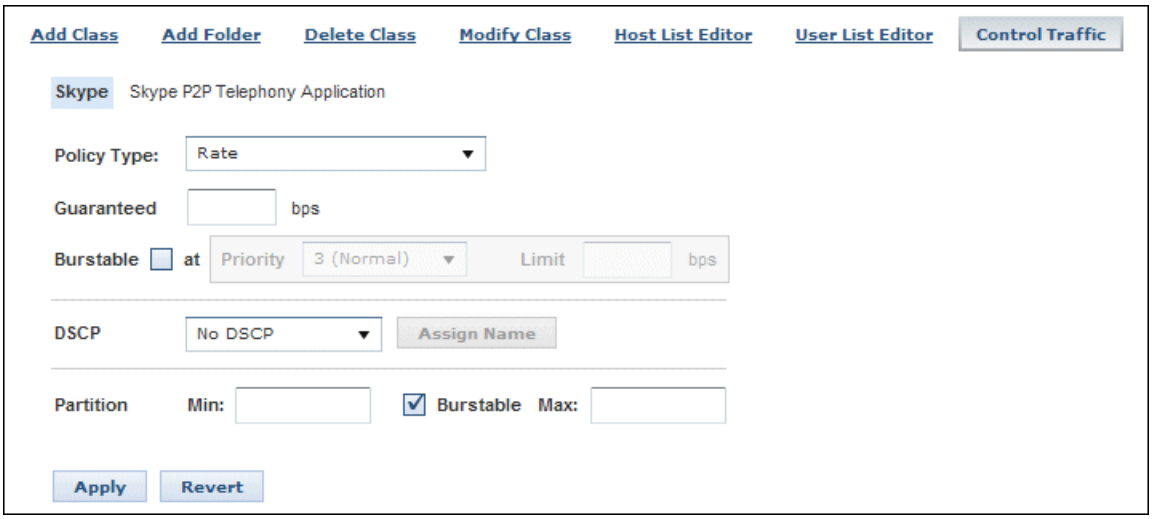

6. From the **Policy Type** drop-down menu, select **Discard**.

7. Click **Apply**.

Additional Information

- Be careful when applying a discard policy, because it may have unexpected results. For example, if you configure the /Inbound/Default class to discard packets, PacketShaper will discard packets for all classes in the Inbound tree that do not have their own policies.
- If you get an error message, you can look at the operation log.

Set an Ignore Policy in the Sky UI

Use an *ignore policy* to treat a traffic type as pass-through traffic so it will not be counted as part of the link traffic under management.

If PacketShaper sees traffic that is *not* destined for the router of the link that it is managing, this traffic should be ignored by PacketShaper, and should not be counted against the capacity of the link. For example, traffic to and from an intranet server will never consume router bandwidth, so it should be ignored.

To apply an ignore policy:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

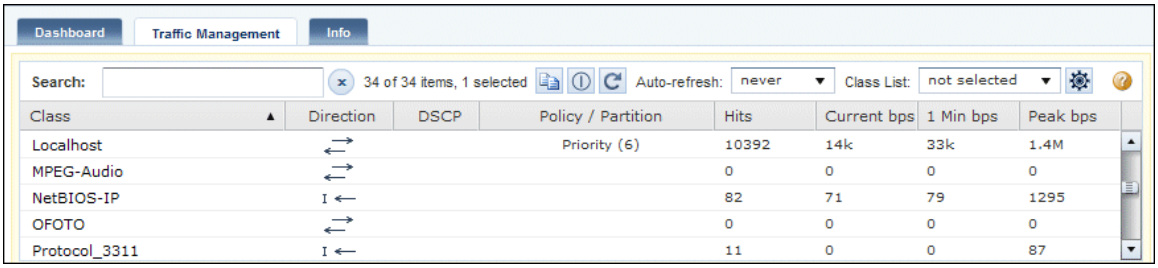

2. Select the class(es) for which you want to assign an ignore policy.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Control Traffic**. **show screen...**

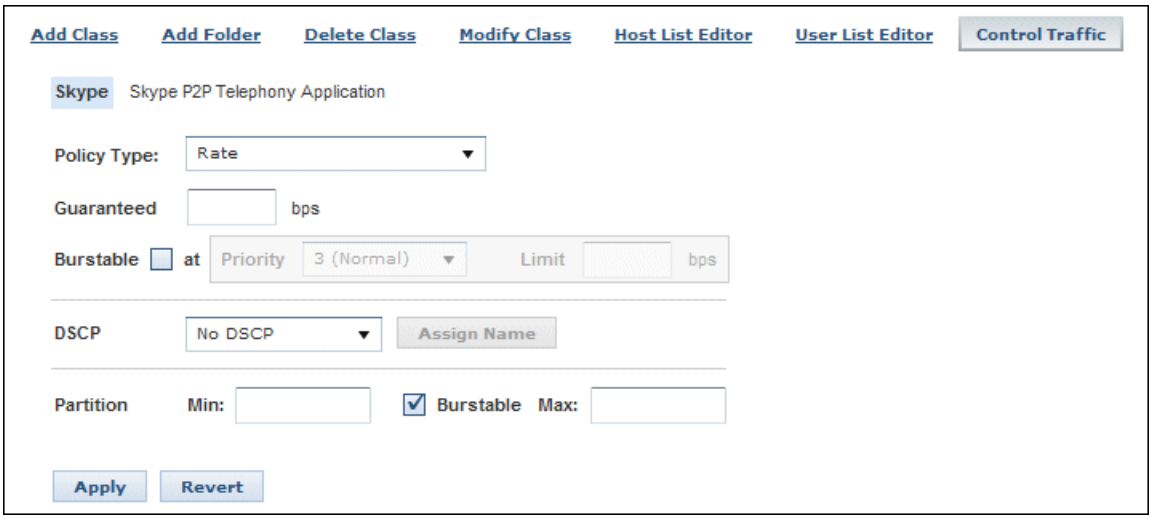

6. From the **Policy Type** drop-down menu, select **Ignore**.

7. Click **Apply**.

Note: If you get an error message, you can look at the operation log.

VoIP Optimization

The Sky UI offers a screen dedicated to optimizing VoIP traffic and recommends appropriate policy settings based on the codec. Sky will recommend the per-flow (that is, per-call) settings for the rate policy. You can also create a partition to guarantee (or limit) bandwidth for all calls of this type.

> **Note:** The VoIP Optimization tab is available when the selected class is based on the RTP-I service and has a codec assigned to it. See Create a VoIP Class in the Sky UI.

To optimize a VoIP class:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

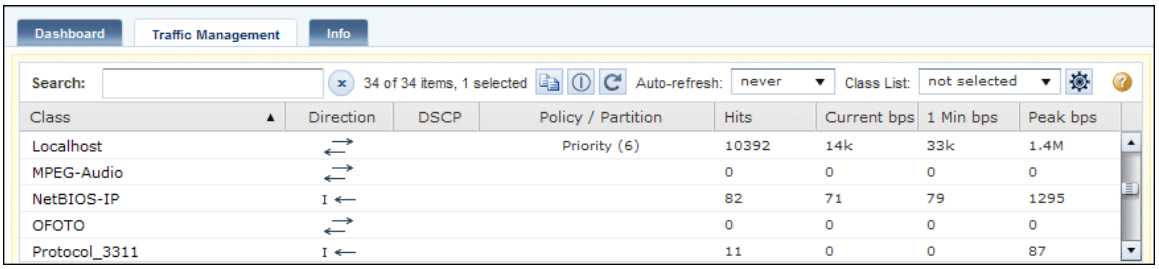

2. In the class tree, select the VoIP class. Only one class can be selected.

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. Select the **VoIP Optimization** tab. The recommended settings appear next to each field in the *Individual Calls* section. **show screen...**

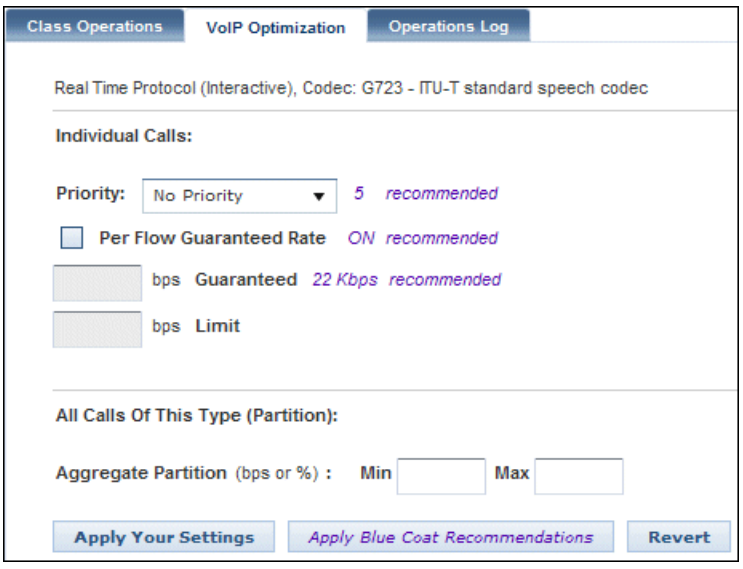

5. Specify the settings for individual calls. To use the recommended settings suggested for each field, click **Apply Symantec Recommendations**. These settings will be entered into the form.
or
Enter your own values:

- Select a **Priority** (0 through 7) from the pull-down menu. Note that 0 is the lowest priority and 7 is the highest. The priority level defines how the class will compete with other classes for excess bandwidth.
- To guarantee a specific rate for each flow (call), enable the **Per Flow Guaranteed Rate** checkbox and enter the rate in the **Guaranteed** field.
- (Optional) To cap the excess rate available to the selected class, enter a value in the **Limit** field.
- Click **Apply Your Settings**.

6. You may also want to guarantee bandwidth for all calls of this type, by creating a partition. Enter the minimum partition size in the **Min** field , expressed in bits per second (bps) or as a percentage of the parent partition's size (%). If the size is entered in bits per second, you can follow the integer by a “k” (thousands), “M” (millions), or "G" (billions). For example, if you want to support 10 concurrent calls at 25 Kbps each, you would enter about 250 Kbps for the partition minimum.

7. (Optional) Select the **Burstable** checkbox to allow the partition to use available excess bandwidth if needed. When this box is unchecked, the partition has a fixed size; when it is not using its reserved bandwidth, that bandwidth is available to other traffic.

8. (Optional) If you selected **Burstable** in step 7, you can enter a value in the **Max** field to cap the amount of excess bandwidth for the partition. If you don't enter a value, the partition size has no maximum limit. The value can be specified in bits per second or as a percentage of the parent partition's size.

9. Click **Apply Your Settings**.

> **Note:** If you get an error message, you can look at the operation log.

DSCP Marking

The Differentiated Services (Diffserv) standard outlines an approach for giving precedence to specific traffic types, providing a scalable solution to managing traffic. You can change the precedence or service level of an application by creating policies to mark the traffic in a particular class with a specific diffserv code point (DSCP) value. The network router can then apply differentiated grades of service to packets based on the DSCP marker in the IP packet header.

You can set the DSCP value for classes that have either rate or priority policies.

To set the DSCP value for one or more traffic classes:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

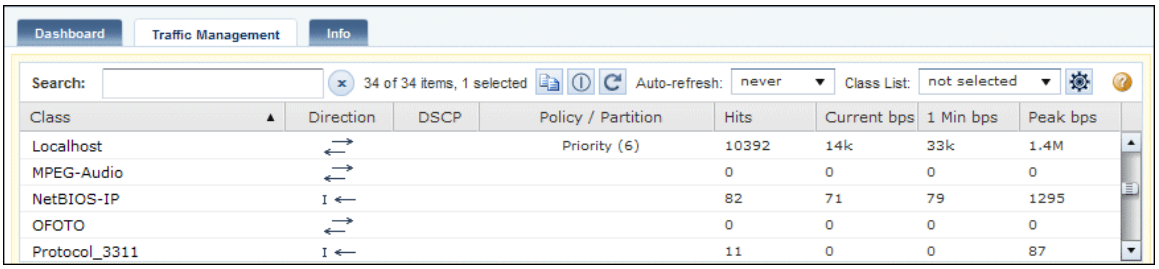

2. Select the class(es) for which you want to assign the same DSCP value.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. Click **Control Traffic**. **show screen...**

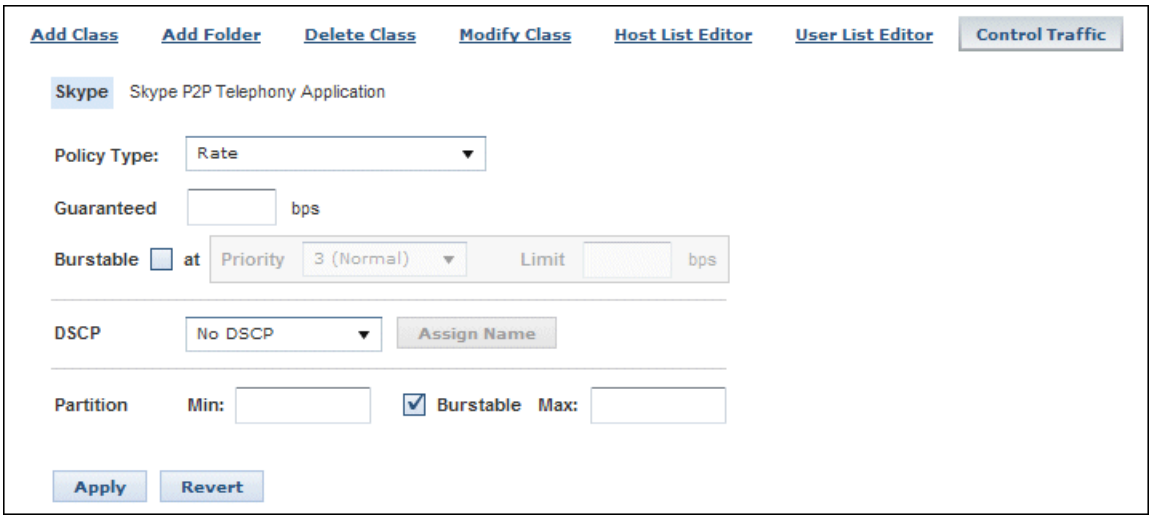

6. For the **Policy Type**, select either Rate or Priority and enter the desired settings.

7. Select a value from the **DSCP** list. **show screen...**

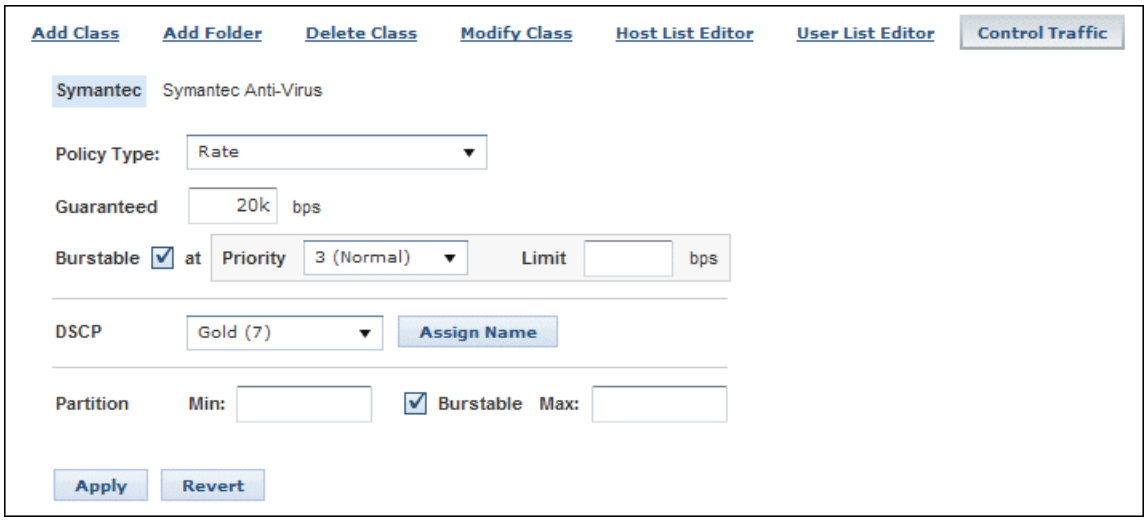

8. If you want to associate a description with the value, click **Assign Name**, enter the name, and click **Done**. DSCP names can be up to 31 characters long; all special characters, including spaces, are allowed.

9. Click **Apply**.

Additional Information

- To change the DSCP name, modify the policy for a class that uses that DSCP name. You can then click **Assign Name**, enter the new name, and click **Done**.
- To remove a DSCP name, modify the policy for a class that uses that DSCP name. Click **Assign Name**, select and delete the current name, and click **Done**.

Delete a Partition in the Sky UI

To delete a partition for a traffic class:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane.**show screen...**

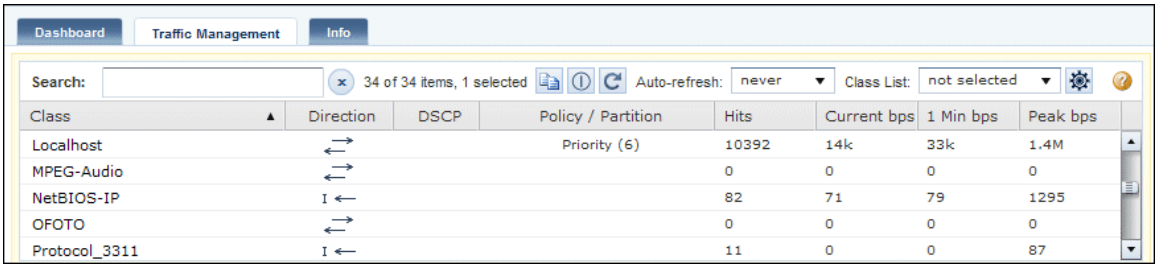

2. Select the class for which you want to delete a partition.

3. Click the **Policy Manager** icon.

4. Click the **Class Operations** tab.

5. In Policy Manager, click **Control Traffic**. **show screen...**

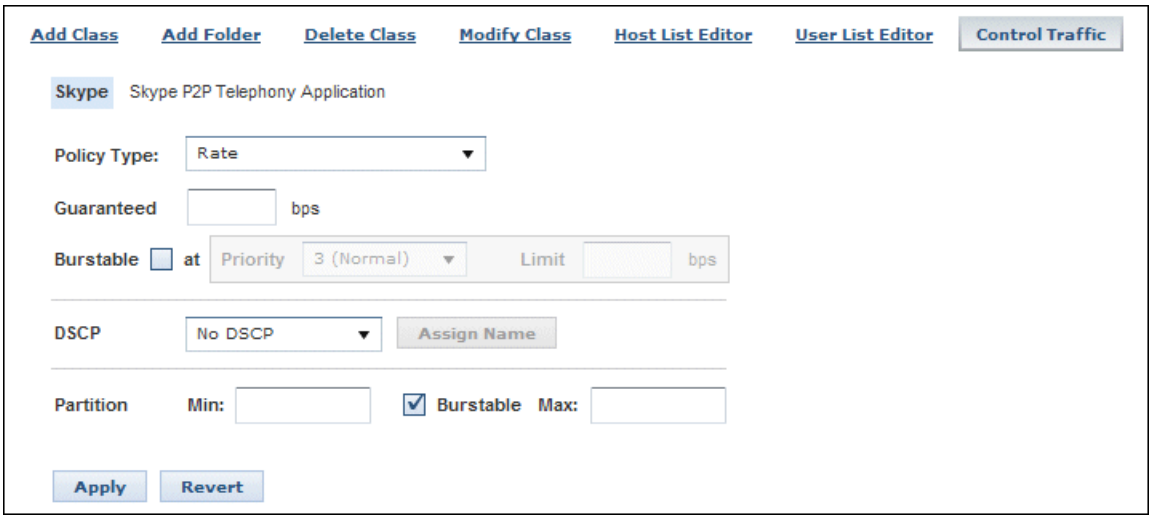

6. For the partition **Min**, delete the value.

7. For the partition **Max**, delete the value.

8. Click **Apply**.

Remove a Policy in the Sky UI

To remove a policy from one or more traffic classes:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

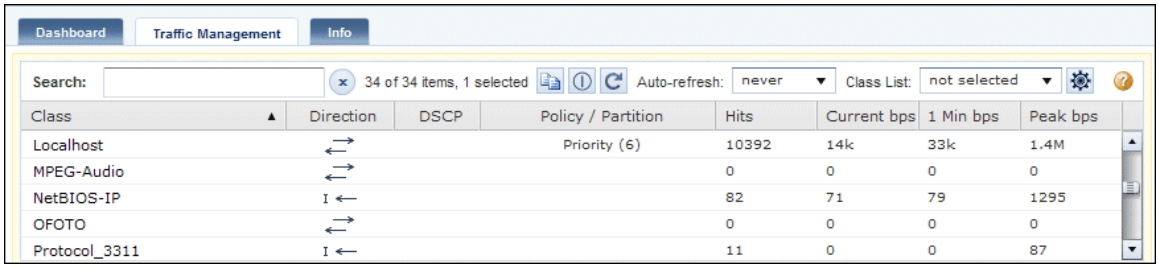

2. In the class tree, select the class(es) for which you want to remove a policy.

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. In Policy Manager, click **Control Traffic**. **show screen...**

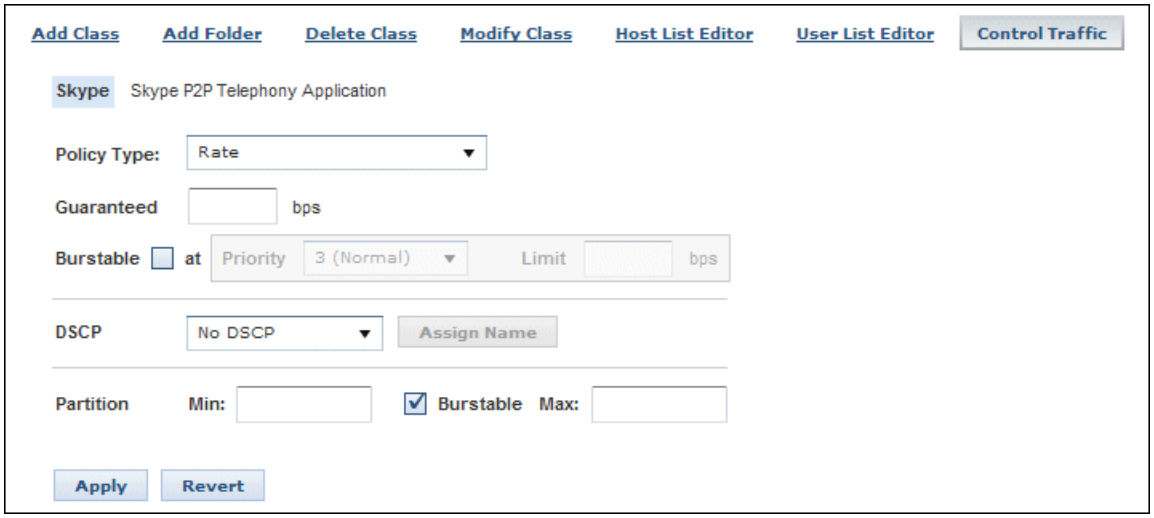

5. For **Policy Type**, select **No Policy**.

6. Click **Apply**.

Modify a Policy/Partition in the Sky UI

To modify traffic ccontrol settings for one or more traffic classes:

1. Click the **Traffic Management** tab. The class tree appears in the upper pane. **show screen...**

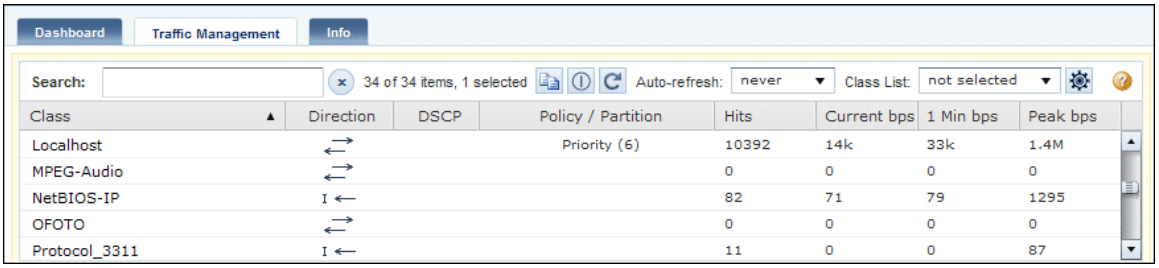

2. In the class tree, select the class(es) for which you want to modify a policy.

3. Click the **Policy Manager** icon. The Policy Manager appears in the lower pane.

4. In Policy Manager, click **Control Traffic**. **show screen...**

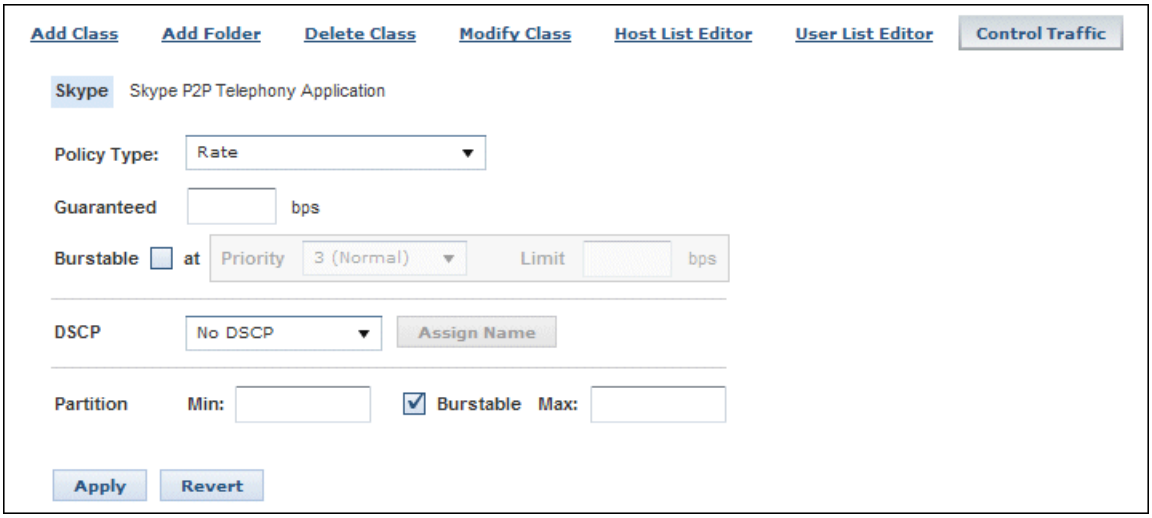

5. Modify the parameters.

6. Click **Apply**.

> **Note:** If you get an error message, you can look at the operation log.

Sky UI Graphs & Reports Tasks

For viewing graphs and reports of your network traffic, the Sky UI offers a Dashboard and a Reports module. The Dashboard gives you an instant picture of what's happening on your network right now, while the Reports module lets you graph specific traffic classes in real time or over a range of time in the past. Any of these graphs can be saved and printed.

View the Sky UI Dashboard

Overview

The Sky UI Dashboard tab displays several real-time graphs that allow you to get a picture of the type of traffic running on your network as well as the current link utilization. The dashboard displays the following graphs:

Top Ten Classes pie charts –Shows the relative portions of bandwidth allocated to the ten most active classes on the network and aggregates all other traffic into a slice called *All others*. The pie chart on the left represents Inbound traffic, while the chart on the right represents Outbound. When mousing over a slice, a tooltip displays the class' average bandwidth usage in bits per second and its percentage of the total bandwidth on the link. Data is refreshed every 10 seconds, showing real-time information for the most recent one-minute interval.

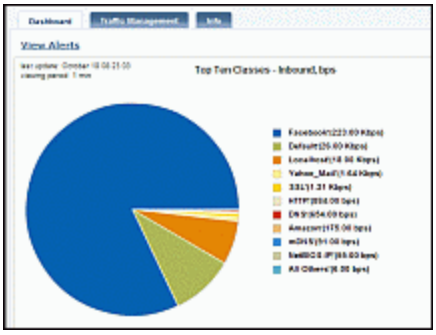

Top Ten Hosts pie charts : Click the **Hosts** button to see the names of the ten users or IPv4/IPv6 addresses consuming the most bandwidth. The pie chart on the left represents the top ten *inside* hosts (relative to the PacketShaper) and the chart on the right represents the top ten *outside* hosts. Note that the pie slices do not represent a portion of the total link, but rather the portion of the top ten host total. When mousing over a slice, a tooltip displays the host's average bandwidth usage in bits per second. Data is refreshed every 10 seconds, showing real-time information for the most recent one-minute interval. User names are included if BCAAA is installed, enabled, and configured, and PacketShaper is able to communicate with the BCAAA server. See Configuring PacketShaper as a BCAAA Client.

> **Note:** The colors in the Inbound (or Inside) and Outbound (or Outside) pie charts correspond to the same classes, users, or IP addresses. For example, if a class appears in the *Top Ten Classes - Inbound* pie chart, the same class will have the same color in the *Top Ten Classes - Outbound* chart.

Link Utilization line graphs – Shows the Inbound and Outbound link's bandwidth usage in bits per second. The graphs are updated every second. The graph initially shows utilization data for the last three minutes; after a period of time the x-axis scale converts to displaying the last 15 minutes of data.

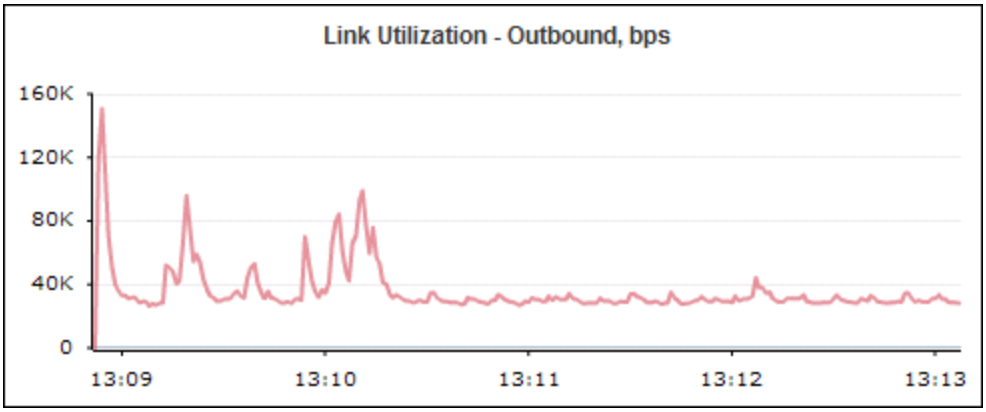

Adjusting the Graph Sizes

Controls are available for resizing the areas allotted to each graph, allowing you to enlarge a particular graph or pair of graphs.

1. Place the mouse pointer on the horizontal or vertical slider icon (for example,). When you have it positioned properly, the pointer will change to double arrows.

2. To adjust the size of the panes, drag the slider between panes. For example, to increase the size of the pie charts, drag the slider down.

3. To adjust the size of each graph, drag the slider between the Inbound and Outbound graphs. For example, dragging to the right increases the size of the Inbound graph while decreasing the size of the Outbound graph.

Save the Sky UI Dashboard

You can save the dashboard in PDF, JPEG, or PNG formats.

To save the dashboard:

1. Click the **Dashboard** tab.

2. Click the **Save Page** icon in the top-right corner of the window.

3. Choose the desired format:

Save Page as PDF – Select this format if you want to print the page directly.
Save Page as JPEG – Select this format if you want to incorporate the page in a report (such as a Word file).
Save Page as PNG – An alternative graphics format to JPEG.

4. In the *Select location* dialog box, browse to the path where you want to save the file, and modify the filename if desired. The default filename is PSSkyPage.

5. Click **Save**.

Additional Information

- To print the PDF, double-click the file and issue the **File > Print** command from Adobe Acrobat.
- The JPEG and PNG graphic files can be imported into reports or emails.

Monitor Traffic in the Sky UI

Go to the **Traffic Management** tab in the Sky UI to see the list of traffic classes and real-time statistics. **show screen...**

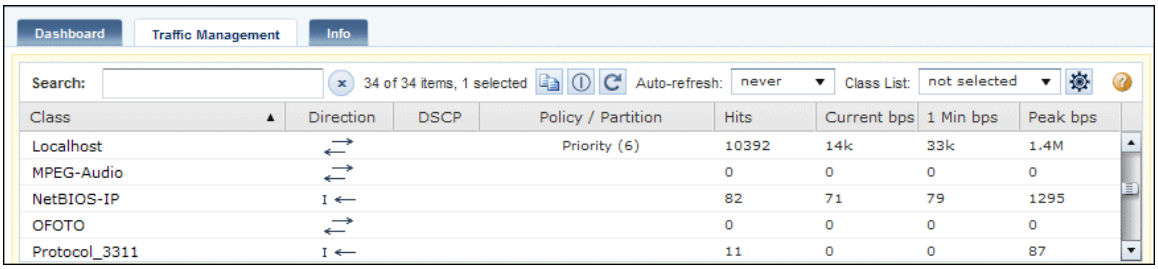

The traffic class window pane lists the following information for each traffic class:

Column	Description
Class	The name of the traffic class.
Direction	Indicates direction of traffic the class is monitoring: ■ Both directions ■ Inbound only ■ Outbound only **Note:** The Direction column appears when the Inbound & Outbound Combined setting is enabled (as it is by default).
DSCP	The configured DiffServ Code Point (DSCP) policy for the traffic class. The DSCP can be a single value, a range of values, or a name (if one has been assigned to a DSCP value).
Policy	Shows the policy (if any) applied to the class. For a priority policy, the defined priority is shown. For a rate policy, the guaranteed rate setting, cap, and burstable priority level are shown.
Partition	The minimum and maximum amount of bandwidth allocated for a class subtree. *Uncommitted* is the difference between the current link speed and the sum of all committed partition minimums. *Uncommitted* indicates that the partition size is dynamically determined at run time and the minimum allocation is whatever is not committed to other partitions. *None* indicates that the partition is burstable without a limit. If the min and max values are the same, the partition is fixed and non-burstable. If a percentage appears after the partition size, the partition was specified as a percentage of the parent partition's size. Example: *450k (10%)* for the Inbound/HTTP class indicates that HTTP's partition size is 10% of the Inbound partition's size, resulting in a partition size of 450 Kbps (assuming the link size is 4.5 Mbps).

Column	Description
Hits	Number of times flows match the traffic class. Class hits occur only at the beginning of a flow or session. Classes created by traffic discovery do not show the first few connections for each traffic type. These initial hits are counted in the Default class. **Note:** In a configuration in which the PacketShaper unit sits between subnets and a router, traffic between two inside hosts will pass through the unit twice: once toward the router, then again back from the router to the other inside host. This results in double hit counts on a SameSide class. For example, four ping packets from one inside host to another would create eight class hits, as opposed to four with a standard setup (inside host to outside host).
Current bps	The current rate values in bits per second
1 Min bps	Last minute's average rate
Peak bps	The current peak values **Note:** Partitions have a 6% tolerance. Therefore, it is possible for an inbound traffic class to have a peak that exceeds the overall inbound peak.

Display Real-Time Graphs in the Sky UI

The Reports module in the Sky UI includes the capability of graphing real-time data of one or more traffic classes. With these graphs, you can view the current real-time bandwidth utilization (in bits per second) for up to 10 classes. A line graph displays the rate at the current time (now) and tracks the rates over the last 15 minutes. The graph updates every second to display real-time data.

To display real-time graphs for one or more traffic classes:

1. Click the **Traffic Management** tab. The class tree appears in the top pane. **show screen...**

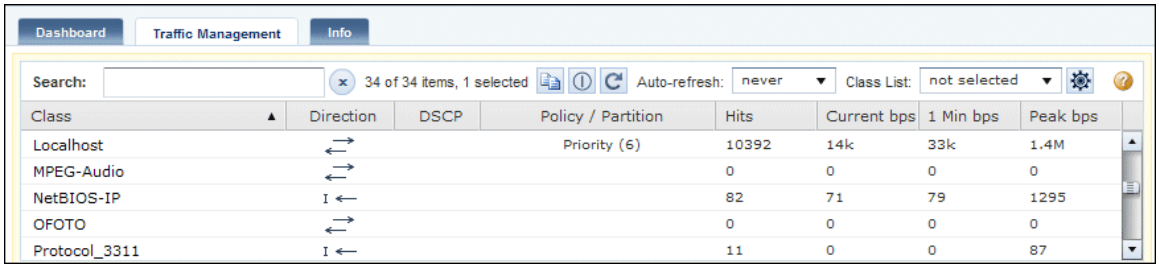

2. Click the **Reports** icon to display the Reports module in the bottom pane.

3. Click the **Real-Time** tab. **show screen...**

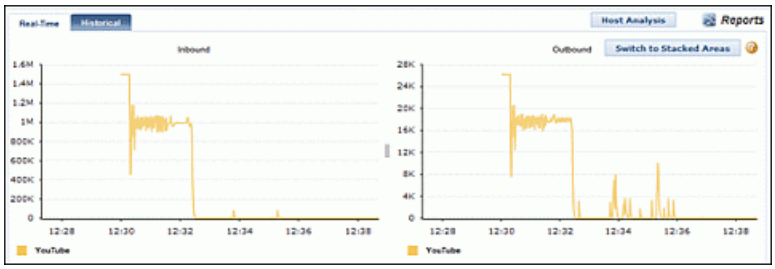

4. Select the class(es) you want to graph, up to 10 classes. Use any of the following techniques for class selection:

- Click anywhere on the class row to select a single class.
- Shift+click to select a range of classes.
- Ctrl+click to select non-adjacent classes.
- Use the **Search** field to locate the class(es) and then select each one.

As you select classes, Sky starts plotting their real-time data, with each class appearing as a different color line. If a class is a combined class, two graphs appear: one for Inbound and one for Outbound. Otherwise, a single Inbound or Outbound graph is displayed.

5. Decide whether you want the data to display as lines or stacked areas. If necessary, click the **Switch to Stacked Areas** button to display a stacked area graph, or the **Switch to Lines** button to display a line graph.

Additional Information

- If you accidentally select the wrong class, either Ctrl+click the class to unselect it or click another class.
- Stacked area graphs are useful for viewing the total volume of traffic.
- Mouse over a data point on a class' line/area to view the real-time rate associated with the point.

Save Graphs in the Sky UI

You can save real-time and historical graphs in PDF, JPEG, or PNG formats.

To save graphs:

1. Display the real-time or historical graph.

2. Resize the graphs to the desired size:

   a. Place the mouse pointer on the horizontal slider icon ( ). When you have it positioned properly, the pointer will change to double arrows.

   b. Drag the slider up or down.

3. Click the **Save Graphs** icon in the Reports module.

4. Choose the desired format:

Save Graphs as PDF – Select this format if you want to print the graphs.
Save Graphs as JPEG – Select this format if you want to incorporate the graphs in a report (such as a Word file).
Save Graphs as PNG – An alternative graphics format to JPEG.

5. In the *Select location* dialog box, browse to the path where you want to save the file, and modify the filename if desired. The default filename is PSSkyGraph.

6. Click **Save**.

Additional Information

- To print the PDF, double-click the file and issue the **File > Print** command from Adobe Acrobat.
- The JPEG and PNG graphic files can be imported into reports or emails.
- The Traffic Management tab has another icon in the top-right corner of the window. This icon saves the entire page (the class tree in the upper pane as well as whichever module is displayed in the lower pane–Reports, Policy Manager, or Settings).

Display Historical Graphs in the Sky UI

The Reports module in the Sky UI includes the capability of graphing historical data for one or more traffic classes. A variety of graph types are available (for example, Utilization with Peaks, Network Efficiency, and TCP Health).

To display historical graphs for one or more traffic classes:

1. Click the **Traffic Management** tab. The class tree appears in the top pane. **show screen...**

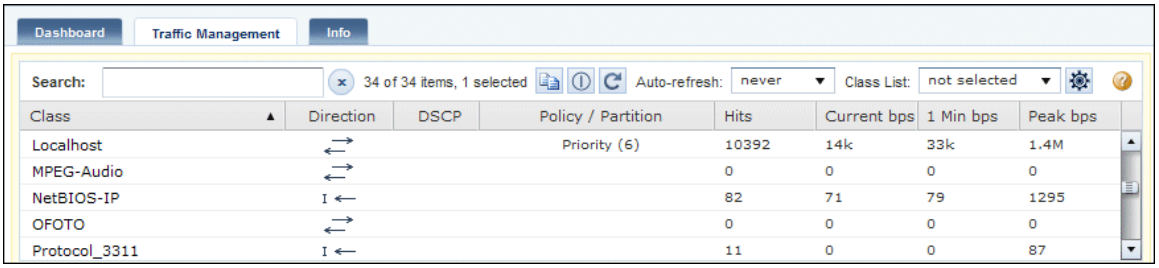

2. Click the **Reports** icon to display the Reports module in the bottom pane.

3. Select the **Historical** tab. **show screen...**

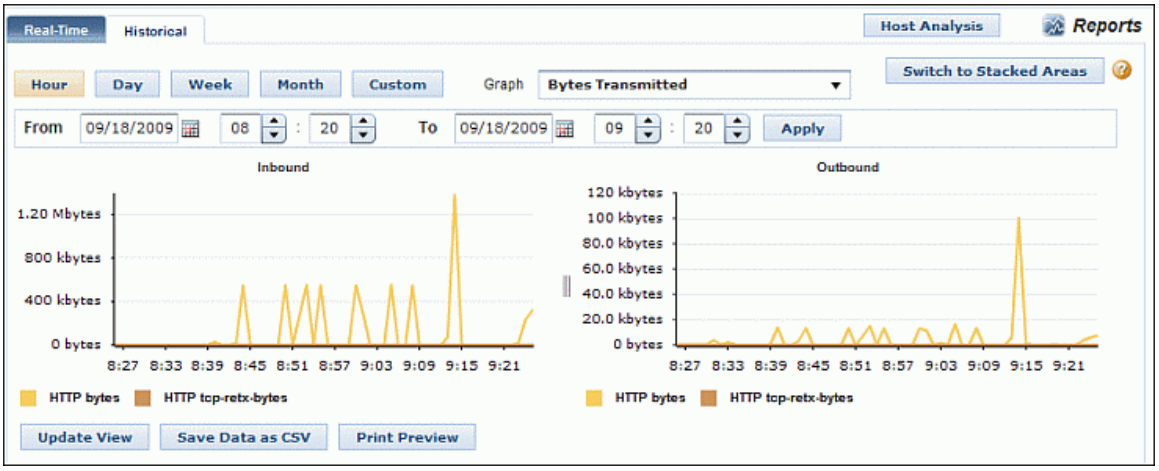

4. Select the graph type from the **Graph** list.

5. Select the time period to graph:

   - Click **Hour**, **Day**, **Week**, or **Month**. These buttons graph the last hour, last day, last week, or last month of data respectively.

   or

   - Click **Custom** and specify a custom time range. In the **From** fields, specify the beginning date/time to graph. In the **To** fields, specify the end date/time to graph.

6. In the class tree above, select the class(es) you want to graph, up to 10 classes. Use any of the following techniques for class selection:

- Click anywhere on the class row to select a single class.
- Shift+click to select a range of classes.
- Ctrl+click to select non-adjacent classes.
- Use the **Search** field to locate the class(es) and then select each one.

7. Decide whether you want the data to display as lines or stacked areas. If necessary, click the **Switch to Stacked Areas** button to display a stacked area graph, or the **Switch to Lines** button to display a line graph.

> **Note:** Stacked area graphs are useful for viewing the total volume of traffic; this style of graph is not appropriate for all graph types.

8. Click **Apply** to view the graph(s) with the specified settings.

Each class/variable combination has a unique color line or area. If a class is a combined class, two graphs appear: one for Inbound and one for Outbound. Otherwise, a single Inbound or Outbound graph is displayed.

Additional Information

- If you accidentally select the wrong class, either Ctrl+click the class to unselect it, or click another class.
- Mouse over a data point on a class' line/area to view the data associated with the point.
- To get a better/larger view of the graphs, you can either increase the size of the Reports pane or click **Print Preview.**
- Use the **Update View** button to refresh the historical graph(s) with current data. This is only applicable to the Hour, Day, Week, and Month graphs.

Print Historical Graphs in the Sky UI

To print an historical graph displayed in the Reports module, you can preview the graph and then issue your browser's Print command.

To print historical graphs:

1. Create an historical graph.

2. Click **Print Preview**.

3. If more than one graph is displayed, a horizontal dividing slider appears between the two graphs. You can drag the slider up or down to adjust the size of the graphs. Drag the slider all the way down to view/print the top graph only; drag the slider all the way up to view/print the bottom graph only.

4. Select **File > Print**.

Save Historical Graph Data in the Sky UI

To save the data associated with an historical graph displayed in the Reports module, you can create a CSV (comma separated value) file. This file can then be imported into Microsoft Excel for further manipulation.

To save historical graph data:

1. Create an historical graph.

2. Click **Save Data as CSV**. The Save dialog box appears. **show screen...**

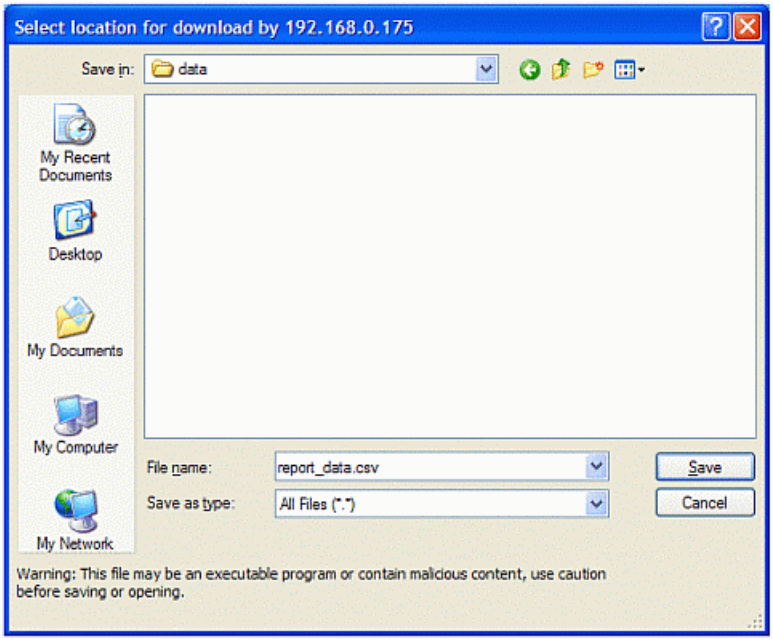

3. Specify a filename and location and click **Save**.

You can now import the CSV text file into Excel.

Reference - Look it up!

Lists and tables of information that can be viewed onscreen or printed out for future reference.

PacketGuide offers the following reference resources:

Command Line Interface (CLI)	Index of all the commands in the PacketShaper 11.9 Command Line Interface (CLI). A PDF version is available for downloading and printing.
Policy/Partition Guidelines	Guidelines for choosing suitable policies for specific types of traffic and for setting appropriate field values
Measurement Variables	List of all common, link, partition, and class measurement variables in PacketShaper 11.9
Applications, Protocols, and Services	List of all applications, protocols, and services for which PacketShaper 11.9 can automatically detect and create classes
URL Categories	A table that lists URL categories, descriptions, and examples.
Advanced UI Graphs	A description of each type of preconfigured graph available in the Advanced UI
Sky UI Graphs	A description of each type of preconfigured graph available in the Sky UI
PacketShaper Port and URL Usage	A list of the ports that PacketShaper uses (and that need to be open on the firewall) and URLs that PacketShaper needs to have access to.
Glossary	Glossary of PacketShaper terms
Tutorials	Links to PacketGuide's multimedia tutorials that demonstrate how to perform common configuration operations
Feature Matrix	A table that lists major PacketShaper features and the version each was introduced or removed.

Policy/Partition Guidelines

Guidelines for choosing suitable policies and partitions for specific types of traffic and for setting appropriate field values

Managing bandwidth allocation for today's traffic diversity is a learned skill. The following table should give you a head start. But keep in mind that these are just guidelines and that your own context heavily influences the types of policies and partitions that are appropriate for your environment.

See Policy Overview and Partition Overview for background information what PacketShaper policies and partitions are and what they do. See Control Strategy Overview for a more broad and introductory approach.

First figure out how to characterize the types of traffic you want to manage, then consult the following tables.

By Traffic Characteristics:

Traffic Characteristics	Traffic Examples	Control Suggestions
Non-TCP traffic	IPX, SNA, AppleTalk	Priority policy with a priority selected from the priority-guidelines table that reflects importance.
Small, time-sensitive, important flows	Telnet, DNS, LDAP, SNMP	Priority policy with a high priority.
Large, bandwidth-hungry, important, non-time-sensitive flows	File transfers, email	Rate policy with 0 guaranteed, burstable, medium priority. Partition to contain the aggregate of all users
Large, bandwidth-hungry, unimportant flows	Games (in a business setting)	Rate policy with 0 guaranteed, burstable, very low priority, and, if desired, a cap of the appropriate per-session Kbps amount that a particular game needs for acceptable performance. Partition to contain the aggregate of all users to less than 5 percent of capacity (or whatever you are willing to devote).
Large, bandwidth-hungry, interactive, time-sensitive flows	SAP	Rate policy with 0 guaranteed, burstable, high priority. Partition to protect the aggregate for all users.
Real-time streaming audio or video flows that need smooth reception and are sensitive to jitter	Streamworks, WindowsMedia	Rate policy with 25 Kbps guaranteed (or the minimum, per-session amount needed for acceptable performance), burstable, medium-to-high priority, and a cap, perhaps 60 Kbps, that prevents sessions from getting bandwidth beyond that which improves reception. Priority policy for the control flows (as opposed to the flows of data) with a high priority. Partition with a size that accommodates the maximum number of allowed users with minimum per-session bandwidth (500 Kbps for 20 users, for example). Admission Control to refuse more than the maximum number of users (or redirect them).
Unimportant, unsanctioned flows that you'd rather not block but want to vigorously contain	Music downloads, URLs of questionable content	Rate policy with 0 guaranteed, burstable, priority 0, and 2 Kbps cap (or the per-session amount you'd like to give). Partition to contain the aggregate of all users to less than 3 percent of capacity (or lower, if desired).

Traffic Characteristics	Traffic Examples	Control Suggestions
Prohibited, unsanctioned flows	Same as previous, but you want to block it	To block UDP traffic, set a discard policy. To redirect web traffic, set a never-admit policy using the **web-redirect** option with the alternate URL. To block web traffic without redirection, set a never-admit policy using the **web-refuse** option. Otherwise, set a never-admit policy with the **refuse** option.
Flows that are *not* destined for the managed link router	Traffic to and from an Intranet server	Ignore policy.
Flows that are part of a contracted amount of bandwidth	ISP customers' bandwidth, dormitory students, office tenants	Static partition for all users' bandwidth. Dynamic partition to allocate each user's or each subnet's bandwidth equitably.
Traffic to and from multiple branches, managed from a main site	SAP, Oracle, Web, and HTTP from Oslo, Milan, and Paris, managed from a main site in London	Should be hub-and-spoke topology. Static partition for each branch with a size equal to branch's WAN capacity. See Control Branch Traffic from a Main Site for instructions.

By Traffic Type:

Application or Protocol	Traffic Details	Control Suggestions
Email, FTP, Telnet, DNS, SNMP, LDAP, and ERP applications		For applications and protocols that are included in the examples in the previous table, consult that table's suggestions.
Web browsing, to a web server (get requests, typically one packet)	Inbound HTTP to an inside server, Outbound HTTP to an outside server	Priority policy with a priority selected from the priority-guidelines table that reflects importance.
Web browsing, from a web server (responses, typically large and graphic laden)	Inbound HTTP from an outside server, Outbound HTTP from an inside server	Rate policy with 0 guaranteed, burstable, a middle priority, and a cap of the appropriate per-session amount that would keep high-capacity web users from usurping the bandwidth for other web users (100 Kbps, for example, if capacity is available). Partition to contain the aggregate of all web users to less than 40 percent of capacity (or whatever you are willing to devote).

Application or Protocol	Traffic Details	Control Suggestions
Voice over IP Voice clients typically use UDP streams. H.323, the industry standard, starts a conversation on one port (H.323), jumps to another port (Q.931), and eventually splits up into a data flow (RTP) and control flow (RTCP). Note: Video over IP is managed similarly, except the bandwidth amounts are greater. For help and details, see Manage Voice and Video Sessions.	Setup traffic (H.323 and Q.931), small	Priority policy at a medium-to-high priority, for example 6.
	RTCP, small and intermittent	Priority policy at a medium-to-high priority, for example 6.
	RTP and other VoIP data protocols, large and data laden	Rate policy with a guaranteed rate (see next comment), burstable at a high priority for the data flows. For determining the guaranteed rate, use the *Monitor Traffic* window to look at several RTP flows. Observe usage characteristics (current, one-minute average, and peak rate). Typically, if a manufacturer claims that its flow requires 8 Kbps, it will actually need 17 to 21 Kbps due to additional overhead and forward error correction. In addition, it is best to overstate the guarantee in rate policies for UDP traffic by 15 to 20 percent. For example: Rate policy with 25 Kbps guaranteed, burstable at 6. ***Note***: Many links do not deliver full-rated bandwidth. For instance, if you have a 128 Kbps link and are running sustained streaming traffic (such as voice), you will probably need to set the link speed 10 percent lower than actual capacity. This lower rate reflects the overhead for framing and routing updates running on the same line. This is critical for slower (sub-T1) links.
Print	Citrix ICA with tag = 3, LPR, TN5250p, TN3287, NetBIOS IP	Rate policy with 0 guaranteed, burstable, priority 2.
Database traffic	Oracle	Try a rate policy with 0 guaranteed, burstable, at priority 6 or a relatively high priority. If latency is a problem, switch to a priority policy at a high priority such as 6.
Citrix application	Citrix-based ICA applications, classified individually or together	Rate policy: 5 to 20 Kbps guaranteed, burstable, priority 4 or 5. If you have a slow link or a larger number of concurrent Citrix users, use less guaranteed rate. Partition on the Citrix parent class of 25 to 50 percent of the link size, burstable, no limit.

Application or Protocol	Traffic Details	Control Suggestions
Encrypted traffic	SSL, IPSec, others	It's impossible to suggest one management strategy for all encrypted traffic because its characteristics and needs vary, just as those for unencrypted traffic vary. If the encrypted traffic is bursty, batch-oriented, and uses TCP (such as FTP over SSL), use a rate policy. If it's not bursty and doesn't use TCP (such as an IPSec tunnel of Telnet traffic), use a priority policy. A partition on IPSec to protect (if it needs protection) and/or to contain (if other traffic's performance is impacted by IPSec's bandwidth demands) is appropriate. A rate policy with guaranteed bandwidth is usually ***not*** appropriate for IPSec traffic, as it may not need the reserved amount or it may need more than what is allocated. A priority policy is a better choice.

Measurement Variables

> **Note:** When you upgrade to a new PacketShaper software version, you may need to reset the measurement data in order to access any new variables introduced in this version.

PacketShaper provides several sets of measurement variables. The Common variables are available for all measurement types. Each measurement type link, partition, and class also has its own set of variables. The values for these variables are collected over the selected interval. The following tables list the variables by type: Common, Link, Partition, and Class.

Table of Common Variables

PacketShaper stores link, partition, and class measurement data for the variables in this table.

Common Variables	Description
avg-bps	Traffic rate, averaged over the interval, in bits per second. For example, to calculate the avg-bps in a one-minute interval the formula is: (bytes in 1 min *8 bits)/60 seconds.
avg-pps	Packets-per-second average rate Formula: pkts / sample-interval-secs
bytes	Number of bytes that passed through the unit during the specified interval
element	ID assigned to the class
guar-rate-allocs	Count of allocation events that were allowed their guaranteed rates during the measurement interval
guar-rate-fails	Count of "denied" admission-control events when there wasn't enough bandwidth to satisfy guaranteed rate demand. This is an indicator of guaranteed rate failures and therefore can be used as a measurement of user satisfaction.
kbytes	Number of bytes transferred during the measurement interval, divided by 1024 and rounded to the nearest kbyte
peak-bps	Peak rate recorded for the link, partition, or class during the specified interval; PacketShaper determines the peak-bps by looking at the rate recorded for the busiest 1/8-second subinterval (that is, the subinterval that had the highest rate) **Note:** You might occasionally see artificially high spikes in the peaks; this happens when the collection and processing of data is paused because higher priority tasks (such as an ME backup) are executing.
peak-guar-rate-flows	Guaranteed rate flow peak count

Common Variables	Description
peak-tcp-conns	Highest number of simultaneous TCP connections recorded for the interval **Note:** The peak number of TCP connections is *not* the same as the peak number of active flows. Some connections may not become active flows due to class policies, such as a class license or never-admit policy.
pkts	Packet count number of packets that passed through the unit during the specified interval
sample-interval-msecs	Time between measurement samples in milliseconds
sample-interval-overruns	Interval overrun count that is, the number of times the measurement engine was too busy to be able to write the data to disk
sample-interval-secs	Time between measurement samples (interval duration) in seconds
tcp-conn-aborts	Number of TCP connections exited as aborted connections; for example, HTTP stop button hits. This number gets incremented when a connection is in progress and an RST is sent. (RSTs are sometimes used as a lazy way to close a connection as an alternative to sending the full orderly shutdown FIN, FIN-ACK, ACK sequence to close a TCP connection.)
tcp-conn-aborts%	Percentage of TCP connections exited as aborted connections Formula: tcp-conn-aborts / tcp-conn-exits
tcp-conn-exits	Number of TCP connections exited. This number gets incremented when PacketShaper sees the end of a flow with an orderly shutdown (FIN, FIN-ACK, ACK) or terminated by an RST.
tcp-conn-inits	Number of TCP connections started. This number gets incremented when PacketShaper sees a SYN packet initiating a new flow.
tcp-conn-self-denies	Number of PacketShaper-initiated TCP connections exited because admission control (refuse) was applied
tcp-conn-self-denies%	Percentage of PacketShaper-initiated TCP connection attempts exited as denied Formula: tcp-conn-self-denies / tcp-conn-exits
tcp-conn-server-ignores	Number of TCP connections exited as ignored connections that is, the server never responded. This number gets incremented when a connection is quarantined (which means the flow limit specified in the flowlimit policy was exceeded) or if either the SYN ACK or SYN ACK ACK are not seen within a one-minute time-out period.
tcp-conn-server-ignores%	Percentage of TCP connections exited as ignored connections Formula: tcp-conn-server-ignores / tcp-conn-exits
tcp-conn-server-refuses	Number of TCP connections refused by the server. This number gets incremented when a SYN is refused using an RST. Generally this happens when a server wants to deny a connection because it's too busy to accept a new connection.
tcp-conn-server-refuses%	Percentage of refused connections Formula: tcp-conn-server-refuses / tcp-conn-exits

Common Variables	Description
tcp-data-pkts	Count of TCP data packets, including retransmits
tcp-early-retx-toss-pkts	Count of TCP tossed retransmissions; PacketShaper controls retransmission time-outs for outbound retransmissions by discarding premature RTO segments.
tcp-early-retx-toss-pkts%	TCP toss rate Formula: tcp-early-retx-toss-pkts / tcp-early-retx-toss-pkts + tcp-retx-pkts
tcp-efficiency%	Percentage of bytes that were *not* retransmitted Formula: (bytes - tcp-retx-bytes) * 100 / bytes
tcp-retx-bytes	Count of TCP retransmitted bytes
tcp-retx-pkts	Count of TCP retransmitted packets, excluding tossed packets
tcp-retx-pkts%	TCP retransmit rate the percentage of packets seen by PacketShaper that were retransmitted Formula: tcp-retx-pkts / tcp-data-pkts

Table of Link Variables

PacketShaper stores link data for the variables in this table.

Link Variables	Description
hostdb-alloc-fails	Number of times the system was unable to allocate resources for host database flows during the measurement interval. *
ipdg-alloc-fails	Number of times the system was unable to allocate resources to classify/control a non-TCP flow during the measurement interval. *
ipv4-avg-bps	Rate of IPv4 traffic, averaged over the interval, in bits per second. This is an exponentially weighted moving average of five seconds of data throughput.
ipv6-avg-bps	Rate of IPv6 traffic, averaged over the interval, in bits per second. This is an exponentially weighted moving average of five seconds of data throughput.
ipv4-bytes	Number of bytes of IPv4 traffic that passed through the unit during the specified interval
ipv6-bytes	Number of bytes of IPv6 traffic that passed through the unit during the specified interval
link-size-bps	Configured link size
noip-avg-bps	Rate of non-IP traffic, averaged over the interval, in bits per second

Link Variables	Description
noip-bytes	Number of bytes of non-IP traffic that passed through the unit during the specified interval
redirect-pkts	Number of packets redirected from the PacketShaper to a cache device (such as the iShared appliance)
returned-pkts	Number of packets that are returned from a cache device (such as the iShared appliance) that need to be de-encapsulated and sent to their destination
pkt-size-histogram	Histogram of number of packets received in different size buckets. The buckets of packet sizes (in bytes) include: [0-63], [64-127], [128-255], [256-511], [512-1023], [1024-1517], [>=1518]
rx-errors	Count of received packets that were dropped due to hardware errors; measures data at the NIC level **Note:** Because this variable measures data at the NIC level, there will be minimal data recorded in the outbound direction when watch mode is enabled (the only outbound traffic will be PacketShaper management).
rx-no-buffers	Count of packets that were dropped due to unavailable buffers; measures data at the NIC level. * **Note:** Because this variable measures data at the NIC level, there will be minimal data recorded in the outbound direction when watch mode is enabled (the only outbound traffic will be PacketShaper management).
rx-pkts-dropped	Count of received packets that were dropped due to: no buffers available (also counted separately in rx-no-buffers variable) no inbound/outbound connections available random dropping This variable measures data at the NIC level. **Note:** Because this variable measures data at the NIC level, there will be minimal data recorded in the outbound direction when watch mode is enabled (the only outbound traffic will be PacketShaper management).
rx-pkts-dropped%	Percentage of received packets that were dropped Formula: rx-pkts-dropped / total-rx-pkts
shaping-mode	Indicates whether traffic shaping is on, off, or has changed during the interval (1=off, 2=changed, 3=on)
tcp-alloc-fails	Number of times the system was unable to allocate resources to classify/control a TCP flow during the measurement interval. *
total-passthru-bytes	Total number of bytes in the total-passthru-pkts
total-passthru-pkts	Total number of packets that were passed through the unit without being controlled
total-rx-bytes	Total number of bytes received, not including bytes that were dropped
total-rx-pkts	Total number of packets received, not including packets that were dropped

Link Variables	Description
total-sameside-pkts	Number of packets that were classified into the SameSide traffic class (a class that measures traffic not destined for the WAN) regardless of whether the packets are dropped or passed through
total-tx-bytes	Total number of bytes transmitted, not including bytes that were dropped
total-tx-pkts	Total number of packets transmitted, not including packets that were dropped
tx-errors	Count of transmitted packets that were dropped due to hardware errors. *
tx-pkts-dropped	Count of transmitted packets that were dropped due to: no route available I/O errors no buffer available
unsolicited-icmp	Detect denial-of-service attacks; detects an ICMP reply when there was no matching request

** Frequent non-zero values indicate the system is running at peak capacity.*

Table of Partition Variables

PacketShaper stores partition data for the variables in this table.

Partition Variables	Description
commitment-failure%	Percentage of one-minute intervals in which packets were dropped when the partition size was less than 6% below the reserved partition size. For example, if packets were dropped in 6 different one-minute intervals during an hour period, the commitment-failure% would be 10 (percent). If no packets were dropped during the interval, the value would be 0. If packets were dropped in a one-minute interval, the value would be 100.
dynamic-cap-count	Number of flows that were given admission control because the dynamic partition already had the maximum number of users
dynamic-live-user	Peak number of live users in a dynamic partition during the measurement interval (*live user* is defined as a user who is actively using the application and won t be usurped if a new user wants a subpartition and none is available)
dynamic-no-partition-count	Number of flows that were denied because the PacketShaper had reached its maximum number of partitions and new subpartitions could not be created
late-drop-bytes	Number of bytes in a partition that were dropped because they had not yet been transmitted after eight seconds or more of delay queuing
late-drop-pkts	Number of packets in a partition that were dropped because they had not yet been transmitted after eight seconds or more of delay queuing

Partition Variables	Description
lowest-fully-satisfied-priority	Lowest priority traffic level that could get all the bandwidth it requested **Note:** Because this variable is sampled multiple times in the interval, the lowest fully satisified priority may be one that occurred only once.
partition-burst-limit-bps	Configured partition burst limit; same as the partition size if the partition is not burstable
partition-over-limit-msecs	Cumulative time during which allocated bandwidth exceeded the partition minimum size, in milliseconds
partition-over-limit-secs	Cumulative time during which allocated bandwidth exceeded the partition minimum size, in seconds
partition-over-limit-time%	Percentage of time that allocated bandwidth exceeded the partition minimum size Formula: partition-over-limit-msecs / sample-interval-msecs
partition-size-bps	Configured partition minimum size
sched-drop-bytes	Number of bytes in a partition that were dropped due to insufficient buffer space
sched-drop-pkts	Number of packets in a partition that were dropped due to insufficient buffer space

Table of Class Variables

PacketShaper stores class measurement data for the variables in this table.

Class Variables	Description
app-availability%	Percentage of time the service is available The app-availability% variable has the following characteristics: ■ It measures the availability when the application is actively in use and bases it upon transactions that are attempted and responded to. The variable does not perform a periodic ping to a server. ■ It reflects the availability of the application based upon the success of connection attempts to a server. Therefore, if no connection attempts are made over a period of time (for example, at 2am), the availability will be reported as 100%, even if the application was temporarily down for maintenance. ■ At the smallest measurement granularity one minute the app-availability% variable reports either 0% or 100%. If all connection attempts within the one-minute period fail, then availability is considered to be 0%. If any connections succeeded or there were no connection attempts, availability is considered to be 100%. Essentially, at one-minute intervals, the app-availability% variable offers a simple flag to indicate whether any connection attempts were successful. ■ At intervals longer than one minute, the percentage is averaged across the interval.
avg-round-trip-time	Average number of milliseconds a packet takes to go from client to server and back again **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
class-hits	Number of times flows match the class; class hits occur only at the beginning of a flow or session
client-flood-block	Number of flows that were blocked due to a client (flow initiator) exceeding the flow limit rate specified in the policy flowlimit command
conn-speed-hist	Speed histogram over well-known speeds. This histogram provides a profile of use, for example, for dial-up modems versus high-speed modems. Use this data to identify which levels are overburdened. These speeds represent the PacketShaper-detected speeds per flow. This data is recorded only for FTP and HTTP connections. All other classes report zeroes.
icmp-attempt-count	Count of attempted ICMP pings
icmp-rtt-count	Count of received icmp-echo_reply messages
icmp-rtt-msecs-sum	Sum of ICMP ping round-trip times in milliseconds
icmp-rtt-msecs-avg	Average ICMP ping round-trip time in milliseconds
icmp-success-rate%	ICMP success rate (icmp-rtt-count/icmp-attempt-count)
drop-bytes%	Percentage of transmitted bytes dropped due to bandwidth saturation Formula: no-bandwidth-drop-bytes / total-bytes-tx
license-overflows	Count of the number of flows that would exceed the licenses

Class Variables	Description
licenses-peak	The largest value of licenses (flows) issued during the interval. Note that the value of the licenses-peak variable is reported as 0 if no licenses were issued during the interval regardless of how many are in use from prior periods.
licenses-total	The value of the class licenses limit at the end of the interval
mos-x10	Mean Opinion Score, a value derived from the R-Factor per *ITU-T Recommendation G.10,* measures VoIP call quality (see rfactor). PacketShaper measures MOS using a scale of 10-50. To convert to a standard MOS score (which uses a scale of 1-5), divide the PacketShaper MOS value by 10.
network-delay-avg	Average number of milliseconds the class' transaction packets spent in transit **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
network-delay-histogram	Histogram of the number of milliseconds the class' transaction packets spent in transit **Note:** This RTM variable collects data only for connection-based TCP traffic classes. Each histogram contains a table of data where the index represents the lower bound of a time range in milliseconds: [25000], [10000], [5000], [2500], [1000], [750], [500], [250], [100], [75], [50], [25], [10], [0]. Table cells contain the number of transactions whose delay time fell within the range represented by the index.
network-delay-median	Median number of milliseconds the class' transaction packets spent in transit (half the class' network delays were shorter, and half were longer) **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
network-delay-msec	The sum of the network delays of all transactions in the specified interval, measured in milliseconds. (Network delay is the time a transaction spends in transit.) This variable is useful for calculating weighted averages across multiple intervals. **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
no-bandwidth-drop-bytes	Number of bytes dropped due to bandwidth saturation
normalized-network-delay-avg	Transaction delay in the network, normalized by transaction size; measured in milliseconds per kilobyte **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
peak-ipdg-conns	Peak number of concurrent non-TCP flows
pkt-exchange-time-avg	Average packet exchange time Formula: pkt-exchange-time / pkt-exchange-time-samples
pkt-exchange-time-msecs	Interval between when a data packet leaves the PacketShaper and its ACK arrives (ACKs to SYN packets are not counted) **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
pkt-exchange-time-samples	Number of pkt-exchange-time samples that were counted **Note:** This RTM variable collects data only for connection-based TCP traffic classes.

Class Variables	Description
policy-hits	Number of times the policy has been enforced
rfactor	Rating Factor. A value derived from metrics such as latency, jitter, and packet loss per *ITU-T Recommendation G.107*, the R-Factor value helps you quickly assess the quality-of-experience for VoIP calls on your network. Typical scores range from 50 (bad) to 90 (excellent).
round-trip-time-msecs	The sum of the round-trip-times (RTT) of all transactions in the specified interval, measured in milliseconds. (RTT is the time a packet takes to go from client to server and back again.) Note that this measurement is taken once per transaction (not once per packet). This variable is useful for calculating weighted averages across multiple intervals. **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
rtp-latency-count	Number of RTP latency samples used to compute latency
rtp-latency-sum	Sum of measured RTP latency
rtp-latency-avg	Average RTP latency in milliseconds; for instance, if the sum of measured latency is 500 ms and the number of latency samples is 25, then the average latency is 20 ms. Formula: rtp-latency-sum / rtp-latency-count **Note:** Measurement data for rtp-latency-avg isn't calculated until you turn on the Enable Latency Calculations system variable.
rtp-expected-pkts	Expected number of RTP packets
rtp-lost-pkts	Number of lost RTP packets
rtp-lost-pkt%	Percentage of expected RTP packets that have been lost
rtp-jitter-count	Number of jitter samples used to compute jitter
rtp-jitter-sum	Sum of measured jitter
rtp-jitter-avg	Average jitter in milliseconds; for instance, if the sum of measured jitter is 100 milliseconds and the number of jitter samples is 25, then the average jitter is 4 milliseconds. Formula: rtp-jitter-sum / rtp-jitter-count
server-flood-block	Number of flows that were blocked due to a server (flow destination) exceeding the flow limit rate specified in the policy flowlimit command
server-delay-avg	Average number of milliseconds required for servers to process the class' transaction requests. The time starts when the server has received all required request packets and ends when the server issues the first packet of the response. **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
server-delay-histogram	Histogram of the number of milliseconds required for servers to process requests for the class' transactions. For example, these measurements could monitor HTTP response for an Internet server, if you set up a host-specific class for the server. **Note:** This RTM variable collects data only for connection-based TCP traffic classes.

Class Variables	Description
server-delay-median	Median number of milliseconds required for servers to process the class' transaction requests (half the class' server delays were shorter, and half were longer) **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
server-delay-msec	The sum of the server delays of all transactions in the specified interval, measured in milliseconds. (Server delay is the time required for servers to process the class' transaction requests.) This variable is useful for calculating weighted averages across multiple intervals. **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
service-level%	Percentage of transactions that satisfied their performance requirement (as defined by the total-delay-threshold variable); the percentage of good or sufficiently speedy transactions.
service-level-errors	Number of one-minute intervals that do not have the required percentage of speedy transactions, as defined by service-level-threshold%
service-level-threshold%	Maximum percentage of slow transactions (as defined by the total-delay-threshold) that each one-minute interval can have and still be considered an acceptable interval **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
slow-transactions	Number of slow transactions (as defined by the total-delay-threshold variable) **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
total-delay-avg	Average number of milliseconds to complete the class' transactions; includes network delay and server delay. **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
total-delay-histogram	Histogram of the number of milliseconds required to complete the class' transactions **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
total-delay-median	Median number of milliseconds required to complete the class' transactions (half the class' total delays were shorter, and half were longer) **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
total-delay-msec	The sum of the delays of all transactions in the specified interval, measured in milliseconds. (Total delay is the time required to complete a transaction; includes network and server delay.) This variable is useful for calculating weighted averages across multiple intervals.
total-delay-threshold	Number of milliseconds that constitutes "too slow" for a total delay (time required for a transaction to complete) **Note:** This RTM variable collects data only for connection-based TCP traffic classes.
total-bytes-tx	Number of bytes transmitted in a traffic class Note that this variable differs from the link variable total-tx-bytes, which measures the total bytes transmitted on a link.
total-trans	Number of transactions (request-response pairs)
trans-bytes	Transaction size for TCP-based applications
trans-bytes-avg	Average number of bytes per transaction Formula: trans-bytes / total-trans
web-response-2XX	Number of HTTP response messages with 2XX success codes for example, **200 OK** and **201 Created**; this variable tracks responses of any class that has web traffic, including synthetic transactions

Class Variables	Description
web-response-3XX	Number of HTTP response messages with 3XX redirection codes for example, **301 Moved** and **302 Found**; this variable tracks responses of any class that has web traffic, including synthetic transactions
web-response-4XX	Number of HTTP response messages with 4XX client error codes for example, **400 Bad Request** and **404 Not found**; this variable tracks responses of any class that has web traffic, including synthetic transactions
web-response-5XX	Number of HTTP response messages with 5XX server error codes for example, **501 Not implemented** and **502 Timed out**; this variable tracks responses of any class that has web traffic, including synthetic transactions

Applications, Protocols, and Services Classified by PacketShaper 11.9

PacketShaper supports a variety of applications, protocols, and services for traffic classification. The traffic discovery process can detect many of them and automatically create traffic classes. You also can create classes manually via the browser interface or the command-line interface, and specify the services and protocols listed in the table shown below.

There are two sources for PacketShaper services:

- Hundreds of services are **built in** to the PacketShaper code.
- **Web applications** are web-based services that are in the Symantec WebPulse database in the cloud. When WebPulse is enabled in PacketShaper, the appliance can classify by any of these 170+ web applications. Symantec is continually adding new web applications to the WebPulse database; see https://sitereview.bluecoat.com/applications.jsp for an up-to-date list. As new applications are added to the WebPulse database, they are available to PacketShaper after it does its automatic daily download of updated map files. For more information about web apps and feature limitations, see " Web Application Classification Overview" on page 131.

Some services in the following table are listed as *service aggregates*. A service aggregate encompasses a combination of services. When you create a traffic class for a service aggregate, multiple matching rules will be created to handle the traffic types represented by the group. For example, DLS (Data Link Switch) is a service aggregate for both the read and write port numbers. When you create a DLS traffic class, two matching rules will be defined to cover both ports. Since a service aggregate represents a combination of matching rules, you cannot specify port numbers for service aggregates (for example, VDOPhone).

Service Groups

The applications and protocols that can be classified by PacketShaper 11.9 are grouped into the following service groups:

ClientServer: Client/Server applications
Collaboration: Software tools for facilitating communication between remote groups
ContentDelivery: Applications that push/deliver content
DatabaseERP: Database and Enterprise Resource Planning applications
DirectoryServices: Applications that provides access to directory services
Email: Email-related protocols
FileServices: File system applications
Games: Online-enabled games
Healthcare: Healthcare-related applications
Internet: Internet infrastructure protocols
Legacy: Legacy protocols
Messaging: Online instant messaging services
Middleware: Applications that function at an intermediate layer
Multimedia: Applications that deliver audio, video and graphics
NetworkMgmt: Protocols used to manage networks
NonIPv4: Non-IP version 4 protocols
P2P: Peer-to-Peer file sharing applications

Print: Network printing protocols
RemoteAccess: Applications to access remote hosts
Routing: Internet routing/switching protocols
Security: Protocols that provide security for network
SocialNetwork: Social network services
VoIP: Voice/video over IP

> **Note:** IPv6 services are in the same service group as their IPv4 counterpart; for example, the FTP and FTP-IPv6 services are both in the Internet service group.

Services Table

Class Name	Description	Source	Version Introduced
ClientServer: Client/Server applications			
FIX	Financial Information eXchange	Built-in	pre-5.0.0
FoldingAtHome	Distributed Computation Screen Saver (foldingathome.stanford.edu)	Built-in	5.0.2
INFOC-RTMS	Attachmate INFOConnect Response Time Monitor System	Built-in	5.0.0
INT-1	Unisys Interactive 1 (2200, ClearPath IX)	Built-in	5.0.0
MATIP	Mapping of Airline Traffic over IP (RFC 2351) service aggregate	Built-in	5.0.0
MATIP-A	MATIP Type A	Built-in	5.0.0
MATIP-B	MATIP Type B	Built-in	5.0.0
NetIQ	NetIQ AppManager	Built-in	5.1.0
OpenConnect-JCP	OpenConnect JCP clients (browser-based access to host applications)	Built-in	pre-5.0.0
PChome	PChome	Web App	9.2.1
PEPGate	Attachmate PEP Gateway (Unisys 2200)	Built-in	5.0.0
Unisys-TCPA	Unisys TCPA (A Series, ClearPath LX, NX)	Built-in	5.0.0
Collaboration: Software tools for facilitating communication between remote groups			
Clearcase	IBM Clearcase version control application	Built-in	8.2.0
CVSpserver	CVS (Concurrent Versions System) pserver	Built-in	5.2.3
Google-Docs	Google On-Line Word Processor	Built-in	9.2.2
GoToMeeting	GoToMeeting	Web App	9.2.1
Hermes	Hermes News Content Manager **Note:** As of v8.3.1, Hermes traffic is no longer autodiscovered by PacketShaper; however, you can create this class manually.	Built-in	8.0.1
LotusNotes	Groupware for collaborative communication	Built-in	pre-5.0.0
MeetingMaker	Meeting Maker	Built-in	5.0.0
Office365	Microsoft Office 365 Applications	Built-in	9.2.1
T.120	Collaboration application	Built-in	pre-5.0.0
Vyew	Vyew	Web App	9.2.1
WebEx	WebEx Real-Time Communications Platform	Web App	5.2.1
ContentDelivery: Applications that push/deliver content			
Alibaba	Chinese e-commerce	Web App	11.2

Class Name	Description	Source	Version Introduced
Amazon	Amazon	Web App	9.2.1
Ameba	Ameba	Web App	9.2.1
Apple-iTunes	Apple iTunes - Music Downloads	Built-in	6.0.2
Apple-Updates	Update Traffic for Apple Products (MAC OS X operating system versions, and Windows versions of iTunes and Safari)	Built-iin	9.2.2
Ariel-419	Infotrieve document delivery system (on port 419)	Built-in	6.0.2
Ariel-422	Infotrieve document delivery system (on port 422)	Built-in	6.0.2
Ask	Ask	Web App	9.2.1
BackWeb	Push technology. Polite BackWeb has an agent on the client to prevent BackWeb background traffic from interfering with other IP network applications.	Built-in	pre-5.0.0
Bing_Search	Bing Search	Web App	9.2.1
Blogger	Blogger	Web App	9.2.1
Chaincast	Chaincast Flexible Content Delivery System	Built-in	5.2.3
Craigslist	Craigslist	Web App	9.2.1
Cyworld_Blog	South Korean website that allows users to discover and post content	Web App	11.2
Disqus	Disqus	Web App	9.2.1
Ebay	Ebay	Web App	9.2.1
EntryPoint	EntryPoint push traffic (formerly PointCast)	Built-in	5.0.0
ESPN	ESPN	Web App	9.2.1
eTrade	eTrade	Web App	9.2.1
FC2_Blog	FC2 Blog	Web App	9.2.1
Flipkart	Indian electronic e-commerce	Web App	11.2
Fox_Sports	Fox Sports	Web App	9.2.1
Google-Books	Google Online Books	Built-in	9.2.2
Google_Search	Google Search	Web App	9.2.1
GoogleEarth	Google Earth (Keyhole) Maps Client	Built-in	8.0.0
Groupon	Groupon	Web App	9.2.1
iCloud	Apple iCloud Service	Built-in	9.2.2
Imgur	Imgur	Web App	9.2.1
iPlayer-Download	BBC iPlayer - Download traffic	Built-in	9.2.2
Kaspersky-Updates	Kaspersky Updates for Internet Security and Anti-Virus Products	Built-in	9.2.2
Kontiki	Kontiki - Content Distribution Network	Built-in	5.2.2
LiveJournal	LiveJournal	Web App	9.2.1
Lycos_Search	Lycos Search	Web App	9.2.1
Marimba	Marimba Castanet push traffic	Built-in	pre-5.0.0
McAfee-ePo	McAfee ePolicy Orchestrator	Built-in	9.2.2
Microsoft-Updates	Update Traffic for Microsoft Products	Built-in	9.2.1
MLB	Major League Baseball	Web App	9.2.1
Napster2	Napster Pay-Per-Use Music **Note:** The original Napster and Napster P2P clients are classified in the Napster service.	Built-in	6.2.0
Napster2-Data	Napster Music File Downloads	Built-in	6.2.0
Napster2-Other	Other Napster Traffic	Built-in	6.2.0
Naver_Blog	South Korean website that allows users to discover and publish posts	Web App	11.2

Class Name	Description	Source	Version Introduced
NBA	National Basketball Association	Web App	9.2.1
NCAA_Mens_ Basketball_ Tournament_2012	Major video and audio streams and bracketing programs for March Madness	Web App	9.2.1
NewsStand	NewsStand-Reader - publication subscription service	Built-in	pre-5.0.0
NFL	National Football League	Web App	9.2.1
Office_365_ Sharepoint	Office 365 Sharepoint	Web App	11.2
Orange_Search	Orange Search	Web App	9.2.1
PayPal	PayPal	Web App	9.2.1
PointCast	See EntryPoint	Built-in	pre-5.0.0
Rakuten	Rakuten	Web App	9.2.1
Symantec_Live_ Update	Symantec Live Update	Web App	9.2.1
Taobao	Chinese website for online shopping similar to eBay and Amazon	Web App	11.2
TrendMicro-Updates	Trend Micro Updates for Internet Security and Anti-Virus Products	Built-in	9.2.2
Tumblr	Tumblr	Web App	9.2.1
Twitpic	Twitpic	Web App	9.2.1
Veetle	Veetle Video Service	Built-in	8.5.1
Webshots	Webshots Desktop (photo screensaver application)	Built-in	pre-5.0.0
Yahoo_Search	Yahoo Search	Web App	9.2.1
Wikipedia	Wikipedia	Web App	9.2.1
WordPress	WordPress	Web App	9.2.1
DatabaseERP: Database and Enterprise Resource Planning applications			
BAAN	Baan enterprise management system	Built-in	pre-5.0.0
FileMaker	FileMaker Pro database service aggregate	Built-in	pre-5.0.0
FileMaker-DB	FileMaker database access	Built-in	pre-5.0.0
FileMaker-R	FileMaker network host response	Built-in	pre-5.0.0
JDENet	J. D. Edwards OneWorld JDENet protocol	Built-in	5.0.0
MSSQL	Microsoft Structured Query Language service aggregate	Built-in	pre-5.0.0
MSSQL-Mon	SQL monitoring traffic	Built-in	pre-5.0.0
MSSQL-Server	SQL server traffic	Built-in	pre-5.0.0
Oracle	Oracle database application service aggregate	Built-in	pre-5.0.0
Oracle-netv1	Oracle SQL*Net v1-based traffic (v6, Oracle7) **Note:** Oracle-netv1 traffic is no longer autodiscovered by PacketShaper; however, you can create this class manually.	Built-in	pre-5.0.0
Oracle-netv2	Oracle SQL*Net v2/Net8-based traffic (Oracle7, 8, 8i, 9i, 10g)	Built-in	pre-5.0.0
Oracle-JVM-SSL	Oracle JVM (IIOP) traffic over SSL	Built-in	pre-5.0.0
Oracle-SSL	Oracle database over SSL	Built-in	pre-5.0.0
OracleClient	Oracle Java client (Webforms)	Built-in	pre-5.0.0
OracleEM	Oracle Enterprise Manager	Built-in	5.2.1
OracleEM1	Oracle Enterprise Manager-1	Built-in	5.2.1
OracleEM2	Oracle Enterprise Manager-2	Built-in	5.2.1
PostgreSQL	PostgreSQL freeware SQL database	Built-in	6.0.2
Progress	Progress database traffic	Built-in	5.0.0

Class Name	Description	Source	Version Introduced
SAP	SAP - Systemanalyse und Programmentwicklung (Services, Applications, and Products in Data Processing)	Built-in	5.0.0
SAP.MCAST.NET	Multicast Service Announcement Protocol: 224.2.127.254	Built-in	5.0.0
DirectoryServices: Applications that provides access to directory services			
CRS	Microsoft Content Replication Service	Built-in	pre-5.0.0
DHCP	Dynamic Host Configuration Protocol service aggregate	Built-in	pre-5.0.0
DHCP-C	DHCP or BootP Client	Built-in	pre-5.0.0
DHCP-S	DHCP or BootP Server	Built-in	pre-5.0.0
DNS	Domain Name Service	Built-in	pre-5.0.0
DNS-IPv6	Domain Name Service over IPv6	Built-in	9.1.1
Finger	Finger User Information Protocol	Built-in	pre-5.0.0
Ident	Identification Protocol	Built-in	pre-5.0.0
Kerberos	Network Authentication Service (ticket granting and checking)	Built-in	pre-5.0.0
LDAP	Lightweight Directory Access Protocol service aggregate	Built-in	pre-5.0.0
LDAP-Clear	LDAP	Built-in	pre-5.0.0
LDAP-Secure	Secure LDAP	Built-in	pre-5.0.0
mDNS	Multicast DNS (Apple Rendezvous)	Built-in	7.0.0
MS-ActiveDir	Microsoft Active Directory	Built-in	8.2.0
MS-ActiveDir-Rep	Active Directory Replication	Built-in	8.2.0
MS-ActiveDir-Bkup	Active Directory Backup	Built-in	8.2.0
MS-ActiveDir-Rstr	Active Directory Restore	Built-in	8.2.0
MS-Replication	Microsoft File Replication	Built-in	
RADIUS	Service agggregate for Remote Authentication Dial-in User Service	Built-in	pre-5.0.0
RADIUS-Acct	RADIUS accounting service	Built-in	pre-5.0.0
RADIUS-Auth	RADIUS authentication service	Built-in	pre-5.0.0
RRP	NSI Registry Registrar Protocol	Built-in	5.1.0
rwho	UNIX remote who (rwho) command; reports current users for all hosts on the local network	Built-in	pre-5.0.0
TACACS	Login host protocol	Built-in	pre-5.0.0
WHOIS	WHOIS service (application that identifies the owner of a domain name)	Built-in	pre-5.0.0
WINS	Windows Internet Name Service	Built-in	pre-5.0.0
Email: Email-related protocols			
AOL_Mail	AOL Mail	Web App	9.2.1
Biff	UNIX new mail notification	Built-in	pre-5.0.0
ccMail	Lotus cc:Mail email application	Built-in	pre-5.0.0
Comcast_Mail	Comcast Mail	Web App	9.2.1
GMX_FreeMail	GMX FreeMail	Web App	9.2.1
Daum_Mail	Email, address book, and calendar services	Web App	11.2
Google_Mail	Google Gmail	Built-in	9.2.2
Groupwise	Novell Groupwise messaging system service aggregate	Built-in	pre-5.0.0
Groupwise-MTA	Novell Groupwise Message Transfer Agent	Built-in	pre-5.0.0
Groupwise-POA	Novell Groupwise Post Office Agent	Built-in	pre-5.0.0
Hotmail	Hotmail	Web App	9.2.1
IMAP	Interactive Mail Access Protocol service aggregate	Built-in	pre-5.0.0
IMAP-Clear	Interactive Mail Access Protocol	Built-in	pre-5.0.0
IMAP-Secure	Secure Interactive Mail Access Protocol	Built-in	pre-5.0.0

Class Name	Description	Source	Version Introduced
Mail.com	Web portal and web-based email service provider owned by the German internet company United Internet	Web App	11.2
MobileMe_Email	MobileMe Email	Web App	9.2.1
MS-Exchange	Microsoft Exchange	Built-in	8.2.0
MS-Exchange-Dir	Exchange Directory	Built-in	8.2.0
MS-Exchange-MAPI	Exchange MAPI	Built-in	8.2.0
MS-Exchange-Admin	Exchange Admin	Built-in	8.2.0
MS-Exchange-IPv6	Microsoft Exchange over IPv6	Built-in	9.1.1
MS-Exchange-Dir-IPv6	Exchange Directory over IPv6	Built-in	9.1.1
MS-Exchange-MAPI-IPv6	Exchange MAPI over IPv6	Built-in	9.1.1
MS-Exchange-Admin-IPv6	Exchange Admin over IPv6	Built-in	9.1.1
Mail.ru_Mail	Mail.ru Mail	Web App	9.2.1
MSSQ	Microsoft Message Queue service aggregate	Built-in	pre-5.0.0
MSSQ-CQ	MSSQ Client Queue	Built-in	pre-5.0.0
MSSQ-IS	MSSQ Information Store	Built-in	pre-5.0.0
MSSQ-Ping	MSSQ Ping Mechanism	Built-in	pre-5.0.0
MSSQ-QMT	MSSQ Queue Manager Traffic	Built-in	pre-5.0.0
MSSQ-SQ	MSSQ Server Queue	Built-in	pre-5.0.0
Nate_Mail	Korean email service	Web App	11.2
Naver_Mail	Email service of Naver, a search portal in South Korea	Web App	11.2
Office_365_Exchange	Microsoft email application	Web App	11.2
OSI	Open System Interconnection (OSI) over TCP (RFC2126), e.g., Microsoft Exchange X.400	Built-in	pre-5.0.0
POP3	Mail reception (Post Office Protocol) service aggregate	Built-in	pre-5.0.0
POP3-Clear	Post Office Protocol for e-mail	Built-in	pre-5.0.0
POP3-Kerberos	Secure Mail Reception (Post Office Protocol with Kerberos)	Built-in	5.0.2
POP3-Secure	Secure Post Office Protocol for e-mail	Built-in	pre-5.0.0
POP3-IPv6	Mail reception (Post Office Protocol) over IPv6 service aggregate	Built-in	9.1.1
POP3-Clear-IPv6	Post Office Protocol over IPv6 for e-mail	Built-in	9.1.1
POP3-Kerberos-IPv6	Secure Mail Reception over IPv6 (Post Office Protocol with Kerberos)	Built-in	9.1.1
POP3-Secure-IPv6	Secure Post Office Protocol over IPv6 for e-mail	Built-in	9.1.1
QQ_Mail	QQ Mail	Web App	9.2.1
Rambler.ru_Mail	Russian email service	Web App	11.2
SMTP	Simple Mail Transport Protocol (mail transmission) service aggregate	Built-in	pre-5.0.0
SMTP-Clear	Mail transmission	Built-in	pre-5.0.0
SMTP-Secure	Secure mail transmission with SSL	Built-in	pre-5.0.0
SMTP-IPv6	Simple Mail Transport Protocol (mail transmission) over IPv6 service aggregate	Built-in	9.1.1
SMTP-Clear-IPv6	Mail transmission over IPv6	Built-in	9.1.1

Class Name	Description	Source	Version Introduced
SMTP-Secure-IPv6	Secure mail transmission with SSL over IPv6	Built-in	9.1.1
UOL_Email	Email service provided by Brazilian ISP, Universo Online	Web App	11.2
Web_de_FreeMail	Web.de FreeMail	Web App	9.2.1
Yahoo_Mail	Yahoo Mail	Web App	9.2.1
FileServices: File system applications			
AFS	Andrew File System service aggregate	Built-in	5.0.2
AFS-FS	Andrew File System file server	Built-in	5.0.2
AFS-VL	Andrew File System volume location database	Built-in	5.0.2
BackupExec-Publish	Backup Exec publish service	Built-in	9.2.2
Box	Box	Web App	9.2.1
CIFS	CIFS - Common Internet File System	Built-in	8.2.0
CIFS-TCP	CIFS over TCP	Built-in	8.2.0
CIFS-NB	CIFS over NetBIOS	Built-in	8.2.0
CIFS-IPv6	CIFS - Common Internet File System over IPv6	Built-in	9.1.1
CIFS-TCP-IPv6	CIFS over TCP over IPv6	Built-in	9.1.1
CIFS-NB-IPv6	CIFS over NetBIOS over IPv6	Built-in	9.1.1
CU-Dev	Fujitsu Device Control (CU-DEV on TCP/IP)	Built-in	pre-5.0.0
CVSup	CVS-Optimized Network File Distribution System	Built-in	5.1.0
Daum_Cloud	Korean cloud-based file storage and file sharing service	Web App	11.2
Digg	Digg	Web App	9.2.1
Dropbox	Dropbox	Web App	9.2.1
Filesanywhere	Remote file storage and file sharing service	Web App	11.2
lockd	NFS file lock daemon	Built-in	pre-5.0.0
MediaFire	File and image hosting website	Web App	11.2
Microsoft-ds	Microsoft Common Internet File System / Server Message Block (SMB) protocol	Built-in	5.0.0
Mobiliti-6-Control	Packeteer Mobiliti v6 - Control	Built-in	8.21
Mobiliti-7-Control	Packeteer Mobiliti v7 - Control	Built-in	8.2.1
Mobiliti-Data	Packeteer Mobiliti - Data	Built-in	8.2.0
MobileMe_iDisk	MobileMe iDisk	Web App	9.2.1
Naver_Ndrive	Korean online file storage/sharing service	Web App	11.2
NDMP	Network Data Management Protocol	Built-in	9.2.2
NDMP-Control	NDMP Control Traffic	Built-in	9.2.2
NDMP-Data	NDMP Data Traffic	Built-in	9.2.2
NetBIOS-IP	NetBIOS over IP service aggregate (NetBIOS is a program that allows applications on different computers to communicate within a LAN)	Built-in	pre-5.0.0
NetBIOS-IP- DGM	NetBIOS Datagram Service	Built-in	pre-5.0.0
NetBIOS-IP-NS	NetBIOS Name Service	Built-in	pre-5.0.0
NetBIOS-IP-SSN	NetBIOS Session Service	Built-in	pre-5.0.0
NFS	UNIX Network File System (both TCP and UDP)	Built-in	pre-5.0.0
NW5-CMD	Netware 5 - Compatibility Mode Drivers service aggregate	Built-in	pre-5.0.0
NW5-CMD-TCP	Netware 5 - Compatibility Mode Drivers over TCP	Built-in	pre-5.0.0
NW5-CMD-UDP	Netware 5 - Compatibility Mode Drivers over UDP	Built-in	pre-5.0.0
NW5-NCP	Netware 5 Core Protocol	Built-in	pre-5.0.0
OpenDrive	OpenDrive	Web App	9.2.1
Pastebin	Pastebin	Web App	9.2.1

Class Name	Description	Source	Version Introduced
Perforce	Perforce Software Configuration Management System	Built-in	8.5.1
rsync	UNIX remote file synchronization protocol	Built-in	5.3
Scribd	Scribd	Web App	9.2.1
SPN	Symantec Protection Network - Online Backup Service	Built-in	9.2.2
SPN-BackupExec	Symantec Online Storage for Backup Exec	Built-in	9.2.2
SPN-OnlineBackup	Symantec Online Backup	Built-in	9.2.2
SunND	Sun Network Disk boot protocol	Built-in	6.1.0
WAFS	Packeteer Wide Area File System Tunnels	Built-in	8.2.0
WDC	Packeteer Wide Dictionary Compression Tunnels - Data	Built-in	8.2.0
WDC-Control	Packeteer Wide Dictionary Compression Tunnels - Control	Built-in	8.2.0
Webde_SmartDrive	WEB.DE Online-Speicher is a German online file storage service	Web App	11.2
Webhard	Korean total data management service that allows data storage, sharing, and transfer	Web App	11.2
Wetransfer	Wetransfer	Web App	9.2.1
Windows_Live_SkyDrive	Windows Live SkyDrive	Web App	9.2.1
YouSendIt	YouSendIt	Web App	9.2.1
Games: Online-enabled games			
AsheronsCall	Asheron's Call network game by Microsoft	Built-in	5.0.2
Battle.net	Blizzard Entertainment online gaming services, including Diablo, Warcraft, and StarCraft	Built-in	5.0.1
Battlefield	BattleField Games	Built-in	8.4.1
CityOfHeroes	City of Heroes - Massively Multi-player Online Role Playing Game	Built-in	8.2.0
Doom	Doom game	Built-in	pre-5.0.0
Everquest	classified as SonyOnline	Built-in	5.2.1
Gamer.com.tw	Gamer.com.tw	Web App	9.2.1
GameSpy	GameSpy Online Gaming	Built-in	8.4.1
Half-Life	Service aggregate for Half-Life network games by Valve and Sierra Studios (including Opposing Forces, CounterStrike, and Team Fortress)	Built-in	5.0.2
Half-Life-TCP	Half-Life chat and server	Built-in	5.0.2
Half-Life-UDP	Half-Life request and game play	Built-in	5.0.2
Kali	Gaming protocol	Built-in	pre-5.0.0
LucasArts	Jedi Knight II: Jedi Outcast	Built-in	5.2.3
MSN-Zone	Microsoft Network Gaming Zone, including Age of Empires	Built-in	5.0.2
MSN-Zone-TCP	Microsoft Network Gaming Zone - TCP traffic	Built-in	5.0.2
MSN-Zone-UDP	Microsoft Network Gaming Zone - UDPtraffic	Built-in	5.0.2
Mythic	Mythic Entertainment games (Dark Age of Camelot)	Built-in	5.2.2
Quake	Quake game service aggregate	Built-in	pre-5.0.0
Quake-A	Quake (port 26000)	Built-in	pre-5.0.0
Quake-B	Quake (port 27500)	Built-in	pre-5.0.0
Quake-II-TCP	Quake II over TCP	Built-in	pre-5.0.0
Quake-II-UDP	Quake II over UDP	Built-in	pre-5.0.0
Quake-III	Quake III Arena	Built-in	5.0.2
Sharefile	Cloud-based file sharing and storage service built for business	Web App	11.2
Slideshare	Oresentation slide hosting service	Web App	11.2
SonyOnline	Games by Sony Online Entertainment (e.g., EverQuest)	Built-in	5.2.1

Class Name	Description	Source	Version Introduced
Sugarsync	Cloud service that enables active synchronization of files across computers and other devices	Web App	11.2
Tribes	Starsiege Tribes network game by Sierra Studios (includes Tribes 2)	Built-in	5.0.2
Unreal	Unreal Tournament PC Game	Built-in	5.0.2
Unreal-Browser	Unreal Tournament - Browser	Built-in	5.0.2
Unreal-Ping	Unreal Tournament - Info Ping	Built-in	5.0.2
Unreal-Play	Unreal Tournament - Game Play	Built-in	5.0.2
Unreal-Status	Unreal Tournament - Status Request	Built-in	5.0.2
WorldofWarCraft	World of Warcraft game	Built-in	8.0
XboxLive	Microsoft Xbox Live Gaming Service	Built-in	8.0
YahooGames	Yahoo! Games	Built-in	5.0.2
Healthcare: Healthcare-related applications			
For details on managing this type of traffic, see Common Performance-Management Strategies for Healthcare Environments.			
DICOM	Digital Imaging and Communications in Medicine	Built-in	5.2.3
HL7	Health Level Seven	Built-in	5.2.3
Internet: Internet infrastructure protocols			
FTP	File Transfer Protocol service aggregate classification - both FTP commands and data	Built-in	pre-5.0.0
FTP-Cmd-Clear	File Transfer Protocol command channel	Built-in	pre-5.0.0
FTP-Cmd-Secure	Secure FTP command channel (SSL)	Built-in	pre-5.0.0
FTP-Data-Clear	FTP data transport channel	Built-in	pre-5.0.0
FTP-Data-Secure	Secure FTP data transfer channel (SSL)	Built-in	pre-5.0.0
FTP-IPv6	File Transfer Protocol over IPv6 service aggregate classification - both FTP commands and data	Built-in	9.1.1
FTP-Cmd-Clear-IPv6	File Transfer Protocol command channel over IPv6	Built-in	9.1.1
FTP-Cmd-Secure-IPv6	Secure FTP command channel (SSL) over IPv6	Built-in	9.1.1
FTP-Data-Clear-IPv6	FTP data transport channel over IPv6	Built-in	9.1.1
FTP-Data-Secure-IPv6	Secure FTP data transfer channel (SSL) over IPv6	Built-in	9.1.1
Gopher	Search application	Built-in	pre-5.0.0
HTTP	Web traffic - Hypertext Transport Protocol	Built-in	pre-5.0.0
HTTP-IPv6	Web traffic over IPv6	Built-in	9.1.1
HTTP-Tunnel	HTTP Tunnel Traffic (traffic that is sent through an HTTP tunnel via an HTTP proxy server on the Internet - perhaps to bypass firewall rules)	Built-in	6.0.0
HTTP-Tunnel-IPv6	HTTP Tunnel Traffic over IPv6	Built-in	9.1.1
ICAP	Internet Content Adaption Protocol	Built-in	
IP	Internet Protocol (not autodiscovered)	Built-in	pre-5.0.0
IPv6	All Internet Protocol traffic over IPv6	Built-in	pre-5.0.0
IPIP	IP-within-IP Encapsulation Protocol	Built-in	5.2.1
NNTP	Usenet NetNews Transfer Protocol	Built-in	pre-5.0.0
NNTP-Clear	Clear Text Usenet newsgroup transmission	Built-in	pre-5.0.0
NNTP-Secure	Secure Usenet newsgroup transmission	Built-in	pre-5.0.0
SOAP-HTTP	Simple Object Access Protocol over HTTP	Built-in	7.0.0

Class Name	Description	Source	Version Introduced
TCP	Transmission Control Protocol - all Internet TCP traffic (not autodiscovered)	Built-in	pre-5.0.0
TCP-IPv6	Transmission Control Protocol over IPv6	Built-in	9.1.1
TFTP	Trivial File Transfer Protocol	Built-in	pre-5.0.0
UDP	User Datagram Protocol - all Internet UDP traffic (not autodiscovered)	Built-in	pre-5.0.0
UDP-IPv6	User Datagram Protocol over IPv6	Built-in	9.1.1
UUCP	Unix-to-Unix Copy Protocol	Built-in	pre-5.0.0
WAP	Wireless Application Protocol	Built-in	8.0.0
XTP	Xpress Transport Protocol	Built-in	8.0.0
Legacy: Legacy protocols			
3PC	Third Party Connect Protocol	Built-in	8.0.0
ACTNET	Active Networks	Built-in	8.0.0
AFP	AppleTalk Filing Protocol (AppleShare IP)	Built-in	pre-5.0.0
Argus	Argus Protocol (System for Distributed Programs)	Built-in	8.0.0
ARIS	Aggregate Route-based IP Switching	Built-in	8.0.0
AX-25	AX.25 Frames	Built-in	8.0.0
BBN-RCC-Monitor	BBN RCC Monitoring	Built-in	8.0.0
BNA	BNA Protocol	Built-in	8.0.0
BR-SAT-MON	Backroom SATNET Monitoring	Built-in	8.0.0
BulkDataXfer	Bulk Data Transfer Protocol	Built-in	8.0.0
CFTP	File Transfer Protocol over IP	Built-in	8.0.0
Chaos	CHAOS Reliable Streaming Protocol	Built-in	8.0.0
CompaqPeer	Compaq Peer Protocol	Built-in	8.0.0
CPHB	Computer Protocol Heart Beat	Built-in	8.0.0
CPNX	Computer Protocol Network Executive	Built-in	8.0.0
CRTP	Combat Radio Transport Protocol	Built-in	8.0.0
CRUDP	Combat Radio User Datagram Protocol	Built-in	8.0.0
DCN-Measure	DCN Measurement Subsystems	Built-in	8.0.0
DDP	Datagram Delivery Protocol	Built-in	8.0.0
DDX	D-II Data Exchange	Built-in	8.0.0
DGP	Dissimilar Gateway Protocol	Built-in	8.0.0
DISTFS	Distributed File System	Built-in	8.0.0
Emcon	Network Emission Control Protocol	Built-in	8.0.0
Encap	Encapsulation Header Protocol	Built-in	8.0.0
EtherIP	Ethernet-within-IP Encapsulation	Built-in	8.0.0
FC	Fibre Channel Protocol	Built-in	8.0.0
FIRE	FIRE Protocol	Built-in	8.0.0
FNAonTCP	FNA on TCP service aggregate	Built-in	pre-5.0.0
FNAonTCP-1	Transport Independent Convergence (FNA on TCP port 492)	Built-in	pre-5.0.0
FNAonTCP-2	Transport Independent Convergence (FNA on TCP port 493)	Built-in	pre-5.0.0
GMTP	GMTP Protocol	Built-in	8.0.0
HostMonitor	Host Monitoring Protocol	Built-in	8.0.0
IATP	Interactive Agent Transfer Protocol	Built-in	8.0.0
IDPR	Inter-Domain Policy Routing Protocol	Built-in	8.0.0
IDPR-CMTP	IDPR Control Message Transport Protocol	Built-in	8.0.0
IDRP	Inter-Domain Routing Protocol	Built-in	8.0.0
IFMP	Ipsilon Flow Management Protocol	Built-in	8.0.0

Class Name	Description	Source	Version Introduced
IL	Internet Link (IL) Transport Protocol	Built-in	8.0.0
INLSP	Integrated Net Layer Security Protocol	Built-in	8.0.0
IPCV	Internet Packet Core Utility	Built-in	8.0.0
IPLT	IPLT Protocol	Built-in	8.0.0
IPPC	Internet Pluribus Packet Core	Built-in	8.0.0
IPX-in-IP	IPX in IP	Built-in	8.0.0
IRTP	Internet Reliable Transaction Protocol	Built-in	8.0.0
ISIS	ISIS over IPv4	Built-in	8.0.0
ISO-TP4	ISO Transport Protocol Class 4	Built-in	8.0.0
Kryptolan	Kryptolan Protocol	Built-in	8.0.0
LARP	Locus Address Resolution Protocol	Built-in	8.0.0
Leaf-1	Leaf-1 Protocol	Built-in	8.0.0
Leaf-2	Leaf-2 Protocol	Built-in	8.0.0
MERIT-INP	MERIT Internodal Protocol	Built-in	8.0.0
MFE-NSP	MFE Network Services Protocol	Built-in	8.0.0
MHRP	Mobile Host Routing Protocol	Built-in	8.0.0
MICP	Mobile Internetworking Control Protocol	Built-in	8.0.0
MPLS-IP	MPLS encapsulation in IP	Built-in	8.0.0
MTP	Multicast Transport Protocol	Built-in	8.0.0
NARP	NBMA Address Resolution Protocol	Built-in	8.0.0
NetBEUI	NetBEUI - Network protocol for PCs	Built-in	pre-5.0.0
NetVoice	Network Voice Protocol	Built-in	8.0.0
NSFNET-IGP	NSFNET Interior Gateway Protocol	Built-in	8.0.0
PGM	PGM Reliable Transport Protocol	Built-in	8.0.0
PIPE	Private IP Encapsulation within IP	Built-in	8.0.0
PktRadioMeasure	Packet Radio Measurement Protocol	Built-in	8.0.0
PNNI	PNNI over IP	Built-in	8.0.0
PRIV-ENCRYPT	Private Encryption Scheme	Built-in	8.0.0
PTP	Performance Transparency Protocol	Built-in	8.0.0
PVP	Packet Video Protocol	Built-in	8.0.0
QNX	QNX Realtime OS Protocol	Built-in	8.0.0
RSVP-E2E-Ignore	Protocol for Aggregation of RSVP	Built-in	8.0.0
RVD	MIT Remote Virtual Disk Protocol	Built-in	8.0.0
SAT-EXPAK	SATNET and Backroom EXPAK	Built-in	8.0.0
SATMON	SATNET Monitoring	Built-in	8.0.0
SCCSP	Semaphore Communications Secure Protocol	Built-in	8.0.0
SCPS	Space Communications Protocol Suite	Built-in	8.0.0
SCTP	Stream Control Transmission Protocol	Built-in	8.0.0
SDRP	Source Demand Routing Protocol	Built-in	8.0.0
Secure-VMTP	Secure Versatile Message Transaction Protocol	Built-in	8.0.0
SEP	Sequential Exchange Protocol	Built-in	8.0.0
SKIP	Simple Key management for Internet Protocol	Built-in	8.0.0
SLP	Service Location Protocol	Built-in	pre-5.0.0
SM	SM Protocol	Built-in	8.0.0
SMP	Simple Message Protocol	Built-in	8.0.0
SNP	Sitara Networks Protocol	Built-in	8.0.0

Class Name	Description	Source	Version Introduced
Sprite-RPC	Sprite RPC Protocol	Built-in	8.0.0
SPS	Secure Packet Shield Protocol	Built-in	8.0.0
SRP	SpectraLink Radio Protocol	Built-in	8.0.0
SSCOPMCE	Service Specific Connection Oriented Protocol in a Multilink and Connectionless Environment	Built-in	8.0.0
STP	Schedule Transfer Protocol	Built-in	8.0.0
TCF	TCF Protocol	Built-in	8.0.0
TLSP	Transport Layer Security Protocol	Built-in	8.0.0
TMUX	Transport Multiplexing Protocol	Built-in	8.0.0
TPPlus	TP++ Transport Protocol	Built-in	8.0.0
Trunk-1	Trunk-1 Protocol	Built-in	8.0.0
Trunk-2	Trunk-2 Protocol	Built-in	8.0.0
TTP	Trivial Transport Protocol	Built-in	8.0.0
UDPLite	Lightweight UDP	Built-in	8.0.0
UTI	Universal Tunnel Interface Protocol	Built-in	8.0.0
VINES	Virtual Integrated Network Service	Built-in	8.0.0
VISA	VISA Protocol	Built-in	8.0.0
VMTP	Versatile Message Transaction Protocol	Built-in	8.0.0
VRRP	Virtual Router Redundancy Protocol	Built-in	8.0.0
WB-EXPAK	Wideband EXPAK Protocol	Built-in	8.0.0
WB-MON	WIDEBAND Monitoring	Built-in	8.0.0
WSN	Wang Span Network	Built-in	8.0.0
xNet	Cross Net Debugger	Built-in	8.0.0
ZeroHop-Proto	0-hop Protocol	Built-in	8.0.0
Messaging: Online instant messaging services			
For details on managing this type of traffic, see Control Instant Messaging.			
AIM_Express	AIM Express	Web App	9.2.1
AOL-AIM-ICQ	AOL 8.0 - AOL Instant Messenger & ICQ service aggregate	Built-in	5.0.1
AOL-IM	AOL - Instant Messenger & ICQ Client-Server	Built-in	5.0.1
ICQ-2000	ICQ - ICQ2000 Client to Client Protocol	Built-in	5.0.2
AOL-IM-Talk	AOL Instant Messaging Point-to-Point Talk	Built-in	5.2.1
AOL-IM-IMAGE	AOL Instant Messaging Point-to-Point Chat	Built-in	5.2.1
AOL-IM-File	AOL Instant Messaging Point-to-Point File Transfer	Built-in	5.2.1
AOL-ISP	AOL 8.0 ISP client traffic	Built-in	6.0.0
AOL-Default	Unknown AOL traffic	Built-in	6.0.0
AOL-iChat	Apple iChat client traffic	Built-in	8.0.0
Google_Talk	Google Talk	Web App	9.2.1
IRC	Internet Relay Chat service aggregate	Built-in	pre-5.0.0
IRC-Chat	Internet Relay Chat - General chat traffic	Built-in	5.2.1
IRC-DCC	Internet Relay Chat - Direct Client-to-Client traffic	Built-in	5.2.1
IRC-Secure	Secure Internet Relay Chat with SSL (port 994)	Built-in	pre-5.0.0
IRC-Servers	Internet Relay Chat - Server-to-Server traffic	Built-in	5.2.1
IRC-194	IRC on port 194	Built-in	pre-5.0.0
IRC-6665	IRC on port 6665 (server to server)	Built-in	pre-5.0.0
IRC-6667	IRC on port 6667 (client to server)	Built-in	pre-5.0.0

Class Name	Description	Source	Version Introduced
Jabber	Jabber/Google-Talk Instant Messaging Application **Note**: This service autodiscover traffic generated by Google Talk and other instant messaging clients that use the Jabber protocol.	Built-in	8.2.0
Lotus-IM	IBM Lotus Instant Messaging **Note:** Some Instant Messenger traffic is classified into pre-existing services: Broadcast Service will classify into RTSP, Meeting Data into T.120, and Meeting Updates into LotusNotes. In order to track and manage Lotus messaging traffic most effectively, you may want to group Lotus-IM, RTSP, T.120, and LotusNotes into a single folder.	Built-in	6.0.2
Lotus-IM-CommC	IBM Lotus Instant Messaging - Tunneled Community Service	Built-in	6.0.2
Lotus-IM-CommS	IBM Lotus Instant Messaging - Community Service	Built-in	6.0.2
Lotus-IM-MtgS	IBM Lotus Instant Messaging - Meeting Service	Built-in	6.0.2
Lotus-IM-SrvrEx	IBM Lotus Instant Messaging - Server Exchange	Built-in	6.0.2
Meebo	Meebo Instant Messaging Service	Built-in	8.3.1
MSN-IM-FileXsfer	File Transfer on MSN/Windows-Live-Messenger	Built-in	8.3.1
MSN-IM-Webcam	Webcam for MSN/Windows-Live-Messenger	Built-in	8.3.1
MSN-Messenger	MSN Messenger Chat Service	Built-in	5.0.1
QQ	QQ Applications	Built-in	8.3.1
QQIM	QQ Instant Messenger	Built-in	8.4.1
QQGames	QQ Games	Built-in	8.4.1
QQWeb	QQ Web	Built-in	8.4.1
Snapchat	Photo messaging application	Web App	11.2
Windows_Live_ Messenger	Windows Live Messenger	Web App	9.2.1
Windows-POPUP	Windows Messenger Delivery (RPC); classifies Windows Messenger spam pop-up windows that were sent out with applications such as DirectAdvertising	Built-in	6.0.2
YahooMsg	Yahoo! Messenger	Built-in	pre-5.0.0
Yahoo_Web_ Messenger	Yahoo Web Messenger	Web App	9.2.1
Middleware: Applications that function at an intermediate layer			
ActiveX	Microsoft object-oriented program technologies and tools	Built-in	pre-5.0.0
BITS	Microsoft Background Intelligent Transfer Service	Built-in	6.2.0
CORBA	CORBA Internet Inter-ORB Protocol (IIOP)	Built-in	pre-5.0.0
DCOM	Microsoft Distributed Component Object Model	Built-in	pre-5.0.0
JavaRMI	Java 1.1.4 TCP Remote Method Invocation service aggregate	Built-in	5.0.0
JavaRMI- Act	Java 1.1.4 TCP Remote Method Invocation - Activate	Built-in	5.0.0
JavaRMI - Call	Java 1.1.4 TCP Remote Method Invocation - Call	Built-in	5.0.0
JavaRMI-Reg	Java 1.1.4 TCP Remote Method Invocation - Registry	Built-in	5.0.0
SmartSockets	Tibco SmartSockets traffic	Built-in	6.0.2
SunRPC	Sun Remote Procedure Calls (UDP) service aggregate	Built-in	pre-5.0.0
SunRPC-Call	Sun Remote Procedure Calls	Built-in	pre-5.0.0
SunRPC-PortMap	Sun Remote Procedure Calls Port Mapper	Built-in	pre-5.0.0

Class Name	Description	Source	Version Introduced
Multimedia: Applications that deliver audio, video and graphics			
For details on managing this type of traffic, see Control Streaming Media.			
4oD-Streaming	Channel4 On Demand - Streaming Video	Built-in	9.2.2
Abacast	Abacast distributed streaming technology	Built-in	7.0.0
Amazon_Instant_Video	Amazon Instant Video	Web App	9.2.1
AcTVila	AcTVila - Japanese Network Television	Built-in	8.4.1
BBC	British public service broadcasting corporation	Web App	11.2
Blip_TV	Media platform for web series content	Web App	11.2
Blockbuster	Blockbuster	Web App	9.2.1
Break	Break	Web App	9.2.1
CBS	CBS broadcasting	Web App	11.2
CNN	Cable News Network	Web App	11.2
Dailymotion	Dailymotion	Web App	9.2.1
DragonFly	DragonFly Storm Video Sharing Service	Built-in	8.5.1
eBaums_World	Website featuring entertainment media such as videos, Flash cartoons, and web games	Web App	11.2
FlashVideo	Streaming Flash Video **Note:** Content providers such as ABC and ESPN use Streaming Flash Video to deliver video content over the Internet.	Built-in	8.2.0
Google-Picasa	Google Online Albums	Built-in	9.2.2
GoogleVideo	Google Video Service	Built-in	8.2.0
Hulu	Hulu Online Video	Built-in	9.2.2
Hulu-Streaming	Hulu - Streaming Video	Built-in	9.2.2
Ifeng	Hong Kong-based Mandarin and Cantonese-language television broadcaster	Web App	11.2
Indiatimes	India-based online news portal	Web App	11.2
iPlayer-Streaming	BBC iPlayer - Streaming video	Built-in	9.2.2
Justintv	Justintv	Web App	9.2.1
Lastfm	Lastfm	Web App	9.2.1
LiveLeak	Video sharing website that lets users post and share videos	Web App	11.2
Megavideo	Megavideo	Web App	9.2.1
Metacafe	Metacafe	Web App	9.2.1
Motion	Motion video (ESPN Motion, etc.) using DIGStream	Built-in	7.0.0
MPEG-4	MPEG-4 Audio/Video Content	Built-in	8.5.1
MPEG-Audio	Moving Picture Experts Group - Audio Streams	Built-in	pre-5.0.0
MPEG-Video	Moving Picture Experts Group - Video Streams	Built-in	pre-5.0.0
NBC	National Broadcasting Company	Web App	11.2
Netflix	Netflix	Built-in	9.2.1
NetworkTVBoxMedia	Network Television Box Media Traffic	Built-in	8.4.1
NicoNico	NicoNico	Web App	9.2.1
Ofoto	Ofoto image sharing traffic	Built-in	7.1.0
Ogg	Ogg over HTTP	Built-in	7.1.0
Pandora	Pandora music streaming application	Web App	8.3.1

Class Name	Description	Source	Version Introduced
Podcast	Podcast User Agents Traffic - XML/RSS Feeds **Note:** Classifies traffic generated by the following "podcatcher" clients: TVTonic, NewzCrawler, Juice, WinPodder, and RSS_Radio.	Built-in	8.2.0
QuickTime	QuickTime over HTTP	Built-in	5.0.0
RadioNetscape	Radio@Netscape streaming music application, powered by Spinner	Built-in	6.2.0
Real	Service aggregate for Real Networks streaming audio/video application	Built-in	pre-5.0.0
Real-BackChan	Real Networks multicast back-channel traffic	Built-in	pre-5.0.0
Real-Encoder	Real Networks encoder traffic over HTTP or RTSP	Built-in	pre-5.0.0
Real-Multicast	Real Networks real data transport UDP multicast traffic	Built-in	pre-5.0.0
Real-Player	Real Networks Player traffic over HTTP or RTSP	Built-in	pre-5.0.0
Real-RDT-TCP	Real Networks Real Data transport TCP traffic	Built-in	pre-5.0.0
Real-RDT-UDP	Real Networks Real Data transport UDP traffic	Built-in	pre-5.0.0
Real-RTP-TCP	Real Networks Real-time Transport Protocol TCP traffic	Built-in	5.0.0
Real-RTP-UDP	Real Networks Real-time Transport Protocol UDP traffic	Built-in	pre-5.0.0
Real-Web	Real Networks traffic over HTTP or RTSP	Built-in	pre-5.0.0
Reddit	Entertainment, social networking service and news website where registered community members can submit content, such as text posts or direct links	Web App	11.2
RTSP	Real-Time Streaming Protocol	Built-in	pre-5.0.0
Shoutcast	Shoutcast streaming audio	Built-in	pre-5.0.0
Slingbox	Slingbox Streaming Personal Television	Built-in	8.2.0
SmugMug	SmugMug image sharing traffic	Built-in	8.2.0
ST2	Internet Stream protocol, version 2	Built-in	pre-5.0.0
StreamWorks	StreamWorks Audio and Video	Built-in	pre-5.0.0
Stupid_Videos	Viral video website dedicated to humorous, off-the-wall videos	Web App	11.2
Telly	Telly	Web App	9.2.1
Tudou	Tudou	Web App	9.2.1
Twitch_TV	Live streaming video	Web App	11.2
Ustream	Ustream	Web App	9.2.1
Veoh	Internet television company that allows users to find and watch major studio content, independent productions and user-generated material	Web App	11.2
VideoFrame	Citrix VideoFrame service aggregate	Built-in	pre-5.0.0
VideoFrame-TCP	Citrix VideoFrame control (TCP)	Built-in	pre-5.0.0
VideoFrame-UDP	Citrix VideoFrame streaming data (UDP)	Built-in	pre-5.0.0
Vimeo	Vimeo	Web App	9.2.1
WinampStream	Winamp streaming traffic	Built-in	6.0.2
WinMedia	Microsoft Windows Media service aggregate	Built-in	pre-5.0.0
WinMedia-Mcast	Windows Media Streaming over UDP Multicast	Built-in	5.0.0
WinMedia-MSBD	Windows Media Encoder	Built-in	5.0.0
WinMedia-TCP	Microsoft Windows Media (NetShow) Streaming over TCP	Built-in	pre-5.0.0
WinMedia-UDP	Windows Media Streaming over UDP Unicast	Built-in	pre-5.0.0
Youku	Youku	Web App	9.2.1
YouTube	YouTube Video Service	Web App	8.2.0
NetworkMgmt: Protocols used to manage networks			
CAPWAP	Control And Provisioning of Wireless Access Points Protocol	Built-in	11.9.1

Class Name	Description	Source	Version Introduced
CAPWAP-Control	CAPWAP Control Traffic	Built-in	11.9.1
CAPWAP-Data	CAPWAP Data Traffic	Built-in	11.9.1
Day-Time	Day-Time (port 13)	Built-in	5.0.2
Echo	Echo Protocol	Built-in	5.2.1
Echo-IPv6	Echo Protocol over IPv6	Built-in	9.1.1
FlowRecords	Packeteer Proprietary Flow Detail Records	Built-in	7.0.0
ICMP	Internet Control Message Protocol	Built-in	pre-5.0.0
ICMP-IPv6	Internet Control Message Protocol over IPv6	Built-in	9.1.1
iMS	IShared Management System	Built-in	8.2.0
IPComp	IP Payload Compression Protocol (IPComp); a protocol to reduce the size of IP packets	Built-in	6.0.0
NetFlowV5	NetFlow v5	Built-in	7.0.0
NTP	Network Time Protocol service aggregate	Built-in	pre-5.0.0
ProxySG	Symantec ProxySG Appliance Control traffic	Built-in	8.4.1
ProxySG-IPv6	Symantec ProxySG Appliance Control traffic over IPv6	Built-in	9.1.1
ProxySG-ADN	Symantec ProxySG Appliance Tunnel traffic	Built-in	8.4.1
ProxySG-ADN-IPv6	Symantec ProxySG Appliance Tunnel traffic over IPv6	Built-in	9.1.1
ProxySG-Management	ProxySG Management Console Traffic	Built-in	8.3.3
ProxySG-Management-IPv6	ProxySG Management Console Traffic over IPv6	Built-in	9.1.1
RSVP	Resource Reservation Protocol	Built-in	5.0.0
SMS	Microsoft SMS 2.0 - Service Pack 2 or later	Built-in	5.2.1
SMS-Auth	SMS authentication	Built-in	5.2.0
SMS-Chat	SMS remote chat	Built-in	5.2.0
SMS-File	SMS file transfer	Built-in	5.2.0
SMS-RC	SMS remote control	Built-in	5.2.0
SNMP	Simple Network Management Protocol service aggregate	Built-in	pre-5.0.0
SNMP-Mon	SNMP monitoring traffic	Built-in	pre-5.0.0
SNMP-Trap	SNMP event notification	Built-in	pre-5.0.0
STUN	Session Traversal Utilities for NAT	Built-in	8.5.1
Syslog	UNIX system logging	Built-in	pre-5.0.0
TimeServer	Time Server (port 37)	Built-in	5.0.2
WCCP-Control	Web Cache Communication Protocol Control Traffic	Built-in	8.5.1
WebPulse	Symantec WebPulse Service Traffic	Built-in	8.6.1
BC-WebFilter	Symantec Web Filter BUFF Database Updates	Built-in	8.6.1
Category-Map	URL Category Map Downloads	Built-in	8.6.1
WebPulse-Queries	Queries to the Symantec WebPulse Cloud Service	Built-in	8.6.1
NonIPv4: Non-IP version 4 protocols			
AppleTalk	Apple network protocol	Built-in	pre-5.0.0
CiscoDiscovery	Cisco Router Discovery Protocol	Built-in	pre-5.0.0
DECnet	Digital Equipment Corporation network protocol	Built-in	pre-5.0.0
DRP	DECnet Routing Protocol	Built-in	pre-5.0.0
DTP	Dynamic Trunking Protocol	Built-in	7.2.1
FNA	Fujitsu Network Architecture (a variant of SNA)	Built-in	pre-5.0.0
IPX	Novell networking protocol	Built-in	pre-5.0.0

Class Name	Description	Source	Version Introduced
LAT	DEC Printer Support (Local Area Transport)	Built-in	pre-5.0.0
Miscellaneous	Miscellaneous Legacy Protocols	Built-in	pre-5.0.0
MOP-DL	Maintenance Operations Protocol Dump/Load	Built-in	5.1.0
MOP-RC	Maintenance Operations Protocol Remote Console	Built-in	5.1.0
PAgP	Port Aggregation Protocol	Built-in	7.2.1
PPPoE	Point-to-Point Protocol over Ethernet	Built-in	5.2.1
PPPoE-Control	Point-to-Point Protocol over Ethernet - Control	Built-in	5.2.0
PPPoE-Data	Point-to-Point Protocol over Ethernet - Data	Built-in	5.2.0
PVSTP	Cisco Per-VLAN Spanning Tree Plus protocol (PVST+)	Built-in	7.2.1
RARP	Reverse Address Resolution Protocol	Built-in	pre-5.0.0
SNA	IBM Systems Network Architecture protocol	Built-in	pre-5.0.0
SpanningTree	IEEE802.1 Bridge Spanning Tree	Built-in	pre-5.0.0
VLAN-Bridge	VLAN Bridge Protocol	Built-in	7.2.1
VTP	VLAN Trunking Protocol	Built-in	7.2.1

P2P: Peer-to-Peer file sharing applications

For details on controlling this type of traffic, see Control Peer-2-Peer Downloads.

Note: Traffic for TCP handshakes will be classified in the Default class while a Peer-to-Peer application is looking for hosts that are online. Traffic will get classified into the appropriate P2P class once the data downloads begin.

Aimster	Aimster file sharing application service aggregate	Built-in	5.1.0
Aimster-Cmd	Aimster command traffic	Built-in	5.1.0
Aimster-Data	Aimster file transfer traffic (upload/download)	Built-in	5.1.0
Aimster-Init	Aimster initial connection to redirector	Built-in	5.1.0
Ares	Ares P2P Filesharing Application	Built-in	8.2.0
Audiogalaxy	Audiogalaxy Satellite File Sharing Community (including Rhapsody)	Built-in	5.0.2
BitTorrent	BitTorrent Distributed File Sharing System (Including Azureus, BitComet, Flashget, Limewire, and Utorrent)	Built-in	6.0.1
BT-Data	BitTorrent Data	Built-in	8.2.0
BT-Tracker	BitTorrent Tracker **Note:** BitTorrent clients connect to a tracker in order to receive their initial list of peers. Most of the peer exchange traffic will classify into the BT-Tracker service.	Built-in	8.2.0
Blubster	Blubster File Sharing Application	Built-in	5.0.0
DirectConnect	NeoModus Direct Connect File Sharing Community	Built-in	5.1.0
EarthStationV	EarthStation V p2p	Built-in	7.0.0
EarthV-Search	EarthStation V search traffic	Built-in	7.0.0
EarthV-HTTP	EarthStation V over HTTP	Built-in	7.0.0
EarthV-SSL	EarthStationV secure traffic	Built-in	7.0.0
EarthV-PXP	EarthStation V pxp file transfers	Built-in	7.0.0
eDonkey	eDonkey2000 File Sharing Application service aggregate (including eMule, Flashget, and Overnet)	Built-in	5.2.0
eDonkey-TCP	eDonkey2000 data traffic	Built-in	5.2.0
eDonkey-Ping	eDonkey2000 host location ping requests	Built-in	5.2.0
eXeem	eXeem P2P Filesharing Application	Built-in	8.0.0

Class Name	Description	Source	Version Introduced
FileRogue	FileRogue File Sharing Application	Built-in	5.2.3
Filetopia	Filetopia community traffic	Built-in	6.0.2
Furthurnet	Furthur Network Peer to Peer File Sharing	Built-in	5.3.0
Gnutella	File sharing and distribution network service aggregate (a variety of look-alike clients get classified as Gnutella, for example, Acquisition, Limewire, Phex, Mutella, Morpheus, Qtraxmax, Shareaza and Xolox)	Built-in	5.0.0
Gnutella-Cmd	Gnutella command and query traffic	Built-in	5.0.0
Gnutella-Download	Gnutella file transfer from the outside to the inside	Built-in	5.0.0
Gnutella-Init	Gnutella initial connection	Built-in	5.0.0
Gnutella-Upload	Gnutella file transfer from the inside to the outside	Built-in	5.0.0
Groove	Groove Peer-to-Peer Application	Built-in	5.1.0
Hopster	Hopster Proxy Connection	Built-in	8.0.0
Hotline	Hotline File Sharing Community service aggregate	Built-in	5.1.0
Hotline-TCP	Hotline File Sharing Community - TCP	Built-in	5.1.0
Hotline-UDP	Hotline File Sharing Community - UDP	Built-in	5.1.0
iMesh	User-to-User Media Exchange	Built-in	5.0.0
Joost	Joost P2P IPTV Application	Built-in	8.4.1
KaZaA	KaZaA File Sharing Application service aggregate **Note:** MusicCity Morpheus also classifies as KaZaA	Built-in	5.1.0
KaZaA-Cmd	KaZaA command traffic	Built-in	5.1.0
KaZaA-Download	KaZaA file transfer from the outside to the inside	Built-in	5.1.0
KaZaA-Query	KaZaA peer server query	Built-in	5.2.3
KaZaA-Upload	KaZaA file transfer from the inside to the outside	Built-in	5.1.0
Napster	Napster Music Community service aggregate **Note:** This is the original protocol that classifies Napster or Napster look-alike P2P clients. The pay-for-music Napster protocol is classified with Napster2.	Built-in	5.0.0
Napster-Cmd	Napster command traffic	Built-in	5.0.0
Napster-Data	Napster data traffic (upload/download)	Built-in	5.0.0
Napster-Init	Napster initial connection to napster.com	Built-in	5.0.0
Naptser-UDP	Napster UDP traffic (upload/download)	Built-in	5.0.0
Pando	Pando P2P File Sharing Application	Built-in	8.3.1
PeerEnabler	Altnet over KaZaA (v2.5) and PeerEnabler traffic	Built-in	6.0.2
PPLive	PPLive P2P Streaming Video	Built-in	8.2.0
PPStream	PPStream P2P IPTV Application	Built-in	8.5.1
ScourExchange	ScourExchange File Sharing Community.	Built-in	5.0.0
Scour-Web	Scour Exchange - web traffic	Built-in	5.0.0
Scour-CSC	Scour Exchange - CSC protocol	Built-in	5.0.0
Share	Japanese P2P Filesharing Application	Built-in	8.0.0
SoulSeek	SoulSeek P2P Filesharing Application	Built-in	7.1.0
Spotify-Download	Spotify Music Sharing Server Download	Built-in	9.2.2
Spotify-P2P	Spotify Music Sharing P2P	Built-in	9.2.2
Spotify-Update	Spotify Music Sharing Update	Built-in	9.2.2
Thunder	Thunder P2P Traffic	Built-in	8.5.1
Torpark	Torpark Secure Browser Application Protocol	Built-in	8.3.1
Tripnosis	Tripnosis File Sharing Application	Built-in	5.1.0

Class Name	Description	Source	Version Introduced
TVU-Player	TVU Player - P2P TV	Built-in	8.2.0
Warez	Warez - P2P App	Built-in	8.2.0
Winny	Winny P2P Filesharing Application	Built-in	6.0.2
Winny2	Winny P2P Filesharing Application version 2	Built-in	8.0.0
Winnyp	Winnyp P2P Filesharing Application	Built-in	8.5.3
xB Browser	Classifies as Torpark	Built-in	
Print: Network printing protocols			
IPP	Internet Printing Protocol	Built-in	5.0.0
Printer	UNIX line printer spooler (LPR)	Built-in	pre-5.0.0
T38	T.38 - Fax Over IP	Built-in	8.2.0
tn3287	IBM 3270 print traffic (TN 3287 extensions)	Built-in	pre-5.0.0
tn5250p	IBM 5250 print traffic over Telnet	Built-in	pre-5.0.0
RemoteAccess: Applications to access remote hosts			
ATSTCP	Galileo 2915 Terminal and Printer Traffic	Built-in	5.0.0
Attachmate-GW	Attachmate INFOConnect-e-Vantage Gateway	Built-in	5.0.0
Citrix	Citrix connectivity application service aggregate. Enables any type of client to access applications across any type of network connection.	Built-in	5.1.0
Citrix-ICA	Citrix Independent Computer Architecture (ICA)	Built-in	5.1.0
Citrix-SB	Citrix Server Browsing (UDP)	Built-in	pre-5.0.0
CitrixCGP	Citrix Common Gateway Protocol	Built-in	8.2.0
CitrixIMA	Citrix Integrated Management Architecture service aggregate	Built-in	5.0.1
CitrixIMA-CMC	Citrix IMA Management Console	Built-in	5.0.1
CitrixIMA-Svr	Citrix IMA Server to Server	Built-in	5.0.1
GoToMyPC	GoToMyPC HTTP traffic	Built-in	6.0.2
LogMeIn	Software as a service that provides remote connectivity, collaboration, and support solutions	Web App	11.2
pcAnywhere	Remote management collaboration tool service aggregate	Built-in	pre-5.0.0
pcAnywhere-D	pcAnywhere data	Built-in	pre-5.0.0
pcAnywhere-OD	pcAnywhere data (old port)	Built-in	pre-5.0.0
pcAnywhere-OS	pcAnywhere status (old port)	Built-in	pre-5.0.0
pcAnywhere-S	pcAnywhere status	Built-in	pre-5.0.0
Persona	Persoft Persona service aggregate	Built-in	5.0.0
Persona-Clear	Persoft Persona non-encrypted (port 1917)	Built-in	5.0.0
Persona-Secure	Persoft Persona secure SSL	Built-in	5.0.0
radmin	Remote Administrator (remote control software)	Built-in	6.2.0
RDP	Remote Desktop Protocol - Microsoft Windows Terminal Server	Built-in	pre-5.0.0
RemotelyAnywhere	RemotelyAnywhere Remote Desktop Services	Built-in	7.2.1
rexec	UNIX remote execution protocol	Built-in	pre-5.0.0
rlogin	Remote login	Built-in	pre-5.0.0
rsh	UNIX remote shell command	Built-in	5.1.0
SHARESUDP	ALC over UDP for EDS/Shares (TPF)	Built-in	5.0.0
SMTBF	Attachmate Lantern Gateway	Built-in	5.0.0
Telnet	Telnet terminal service aggregate	Built-in	pre-5.0.0
Telnet-Clear	Network terminal protocol	Built-in	pre-5.0.0
Telnet-Secure	Secure Network terminal protocol	Built-in	pre-5.0.0
Telnet-IPv6	Telnet terminal over IPv6 service aggregate	Built-in	9.1.1

Class Name	Description	Source	Version Introduced
Telnet-Clear-IPv6	Network terminal protocol over IPv6	Built-in	9.1.1
Telnet-Secure-IPv6	Secure Network terminal protocol over IPv6	Built-in	9.1.1
Timbuktu	Timbuktu Pro service aggregate (networked remote control)	Built-in	pre-5.0.0
Timbuktu-Ctl	Timbuktu Control Channel	Built-in	pre-5.0.0
Timbuktu-HS	Timbuktu Handshaking	Built-in	pre-5.0.0
Timbuktu-Obs	Timbuktu Observe Channel	Built-in	pre-5.0.0
Timbuktu-Snd	Timbuktu Send Channel	Built-in	pre-5.0.0
Timbuktu-Xch	Timbuktu Exchange Channel	Built-in	pre-5.0.0
tn3270	Telnet for IBM 3270 terminals and 3270 emulation	Built-in	pre-5.0.0
tn5250	IBM 5250 terminal traffic over Telnet	Built-in	pre-5.0.0
VMware-View	VMware View VDI * * Classification for non-SSL connections between the VMware View client/portal and VMware View Connection server	Built-in	8.5.1
VMware-View-Agent	VMware View Agent control traffic that runs on the remote desktop and makes connections to the VMware View connection server	Built-in	9.2.2
VMware-View-Command	VMware View VDI Command and Administration (configuration and login)	Built-in	8.5.1
VMware-View-Desktop	VMware View VDI Offline Desktop (downloads and syncs)	Built-in	8.5.1
VMware-View-PCoIP	PC-over-IP--protocol for remote workstations and desktops	Built-in	9.2.2
VMware-View-Tunnel	VMware View VDI Tunnel (remote desktop stream)	Built-in	8.5.1
VNC	Virtual Network Computing	Built-in	5.0.0
XWindows	X11 Windowing agent (UDP)	Built-in	pre-5.0.0
XWindows-DM	XWindows Display Manager (XDMCP)	Built-in	pre-5.0.0
XWindows-S	XWindows Server	Built-in	pre-5.0.0
Routing: Internet routing/switching protocols			
AURP	AppleTalk Update-based Routing Protocol	Built-in	pre-5.0.0
BGP	Border Gateway Protocol	Built-in	pre-5.0.0
CBT	Core-Based Trees (Multicast Routing Protocol)	Built-in	pre-5.0.0
EGP	Exterior Gateway Protocol (network routing information)	Built-in	pre-5.0.0
EIGRP	Enhanced Interior Gateway Routing Protocol	Built-in	pre-5.0.0
IGMP	Internet Group Management Protocol	Built-in	pre-5.0.0
IGP	Interior Gateway Protocol	Built-in	5.0.0
OSPF	Open Shortest-Path First network routing information	Built-in	pre-5.0.0
PIM	Protocol-Independent Multicast Routing Protocol	Built-in	pre-5.0.0
RIP	Routing Information Protocol (UDP)	Built-in	pre-5.0.0
SSDP	Simple Service Discovery Protocol	Built-in	5.3.0
Security: Protocols that provide security for network			
DLS	SNA over TCP transport - service aggregate classification of Data Link Switch traffic, both read and write port numbers	Built-in	pre-5.0.0
DLS-RPN	Data Link Switch Read Port Number	Built-in	pre-5.0.0
DLS-WPN	Data Link Switch Write Port Number	Built-in	pre-5.0.0
DPA	Microsoft Distributed Password Authentication	Built-in	pre-5.0.0

Class Name	Description	Source	Version Introduced
Google-SafeBrowsing	Safe Browsing Extension for Firefox	Built-iin	9.2.2
GRE	General Routing Encapsulation	Built-in	pre-5.0.0
HSRP	Cisco Hot Standby Router Protocol	Built-in	9.2.2
IPMobility	Minimal encapsulation for IP [RFC 2004]	Built-in	6.1.0
IPSec	IP Security Encapsulation service aggregate	Built-in	pre-5.0.0
IPSec-AH	IPSec Authentication Header	Built-in	pre-5.0.0
IPSec-ESP	IPSec Encapsulating Security Payload	Built-in	pre-5.0.0
IPSec-NAT	UDP Encapsulation of IPSec for NAT traversal	Built-in	
ISAKMP	ISAKMP/IKE key exchange	Built-in	pre-5.0.0
L2TP	Layer 2 Tunneling Protocol for VPN connections (UDP encapsulation)	Built-in	pre-5.0.0
PPTP	Point-to-Point Tunneling Protocol	Built-in	pre-5.0.0
RC5DES	DES (data encryption standard) encryption-cracking application	Built-in	pre-5.0.0
SOCKS	SOCKSv4 and SOCKSv5 proxy protocol	Built-in	pre-5.0.0
SoftEther	Japanese SSL VPN secure traffic	Built-in	7.1.0
SSH	Secure Shell remote login protocol	Built-in	pre-5.0.0
SSH-IPv6	Secure Shell remote login protocol over IPv6	Built-in	9.1.1
SSL	Secure Sockets Layer protocol (secure Web traffic)	Built-in	pre-5.0.0
SSL-IPv6	Secure Sockets Layer protocol over IPv6 (secure Web traffic)	Built-in	9.1.1
SSL-No-Cert	Secure web traffic (Secure Socket Layer) without Authentication	Built-in	8.3.1
SSL-Shell	SSL Secure Shell remote login protocol	Built-in	pre-5.0.0
SWIPE	swIPe: network layer encapsulated encrypted IP protocol	Built-in	6.1.0
Symantec	Symantec Anti-Virus Application Protocol	Built-in	8.3.1
VPN-TCP	TCP Encapulation of IPSEC traffic for VPNs	Built-in	8.3.1

SocialNetwork: Social network services

Note: The Social Networking URL category includes classification of sites that may not yet have a service in PacketShaper. Therefore, creating a Social_Networking category class may be preferable over creating a SocialNetwork service group class.

Bebo	Bebo	Web App	9.2.1
Classmates	Classmates	Web App	9.2.1
Daum_Blog	South Korean blogging website that allows users to discover and publish posts	Web App	11.2
Facebook	Facebook Social Networking Site	Built-in	8.3.1
Facebook_Plugins	Facebook Plugins	Web App	9.2.1
Flickr	Flickr image sharing traffic	Web App	8.3.1
Friendster	Friendster	Web App	9.2.1
GooglePlus	Social networking and identity service	Web App	11.2
GooglePlus_Widgets	GooglePlus Widgets	Web App	9.2.1
hi5	hi5	Web App	9.2.1
Instagram	Online photo-sharing, video-sharing, and social networking service	Web App	11.2
Kaixin	Kaixin	Web App	9.2.1
LinkedIn	LinkedIn	Web App	9.2.1
meinVZ	meinVZ	Web App	9.2.1
Mixi	Mixi	Web App	9.2.1
Mobile01	Chinese 3C social network where users can participate in online auctions and product forums	Web App	11.2

Class Name	Description	Source	Version Introduced
MySpace	MySpace Social Networking Site	Web App	8.3.1
Ning	Ning	Web App	9.2.1
Odnoklassniki	Odnoklassniki	Web App	9.2.1
Orkut	Orkut Social Networking Site	Built-in	8.3.1
Pengyou	Pengyou	Web App	9.2.1
Plurk	Plurk social networking site	Built-in	9.2.2
QQ_Weibo	Chinese 140-character limit micro blogging site	Web App	11.2
Renren	Renren	Web App	9.2.1
Sina	Sina	Web App	9.2.1
studiVZ	studiVZ	Web App	9.2.1
Tuenti	Spain-based, social networking service	Web App	11.2
Twitter	Twitter	Web App	9.2.1
Twitter_Widgets	Twitter Widgets	Web App	9.2.1
Vine	Mobile app owned by Twitter that enables its users to create and post short looping video clips	Web App	11.2
VK	European social network	Web App	11.2
VKontakte	VKontakte	Web App	9.2.1
Weibo	Weibo	Web App	9.2.1
Wer_kennt_wen	Wer-kennt-wen	Web App	9.2.1
Wretch	Wretch	Web App	9.2.1
Xing	Xing	Web App	9.2.1
Yammer	Yammer	Web App	9.2.1
VoIP: Voice/video over IP			
For details on managing this type of traffic, see Manage Voice and Video Sessions.			
CiscoCTI	Cisco Computer Telephony Interface	Built-in	5.2.3
Clarent-CC	Clarent Voice over IP Command Center	Built-in	pre-5.0.0
Clarent-Complex	Clarent complex traffic	Built-in	pre-5.0.0
Clarent-Mgmt	Clarent management traffic	Built-in	pre-5.0.0
Clarent-Voice-S	Clarent voice traffic (simple)	Built-in	pre-5.0.0
CUSeeMe	Video chat services application service aggregate	Built-in	pre-5.0.0
CUSeeMe-AV	Video chat services audio/video	Built-in	pre-5.0.0
CUSeeMe-CC	Video chat services connection control	Built-in	pre-5.0.0
CUSeeMe-CE	Video chat services connection establishment	Built-in	pre-5.0.0
Dialpad	Dialpad Internet Telephone service aggregate	Built-in	5.2.1
Dialpad-Ctrl	Dialpad Internet Telephone - control traffic	Built-in	5.2.0
Dialpad-Stream	Dialpad Internet Telephone - RTP stream	Built-in	5.2.0
H.323	Internet telephony standard service aggregate	Built-in	pre-5.0.0
H.323-GKD	H.323 Gatekeeper Discovery	Built-in	pre-5.0.0
H.323-H.225-RAS	H.323 Call Signaling and RAS	Built-in	8.7.1
H.323-H.245	H.323 call control	Built-in	pre-5.0.0
H.323-Q.931	H.323 call setup	Built-in	pre-5.0.0
H.323-RAS	H.323 Gatekeeper Control (Registration, Admission, and Status)	Built-in	pre-5.0.0
H.323-DTMF	H.323 Dual Tone Multifrequency Relay	Built-in	8.4.1
I-Phone	Vocaltec Internet telephone service	Built-in	pre-5.0.0
MCK-Signaling	MCK Signaling (not autodiscovered)	Built-in	5.0.0

Class Name	Description	Source	Version Introduced
MCK-Voice	MCK Voice (not autodiscovered)	Built-in	5.0.0
Megaco	Media Gateway Control (H.248)	Built-in	6.0.0
Megaco-Text	Media Gateway Control (H.248) Text	Built-in	6.0.0
Megaco-Bin	Media Gateway Control (H.248) Binary	Built-in	6.0.0
MGCP	Media Gateway Control Protocol	Built-in	6.0.0
MGCP-Gateway	Media Gateway Control Protocol Gateway	Built-in	6.0.0
MGCP-CallAgent	Media Gateway Control Protocol CallAgent	Built-in	6.0.0
MGCP-KpAlive	Media Gateway Control Protocol KeepAlive Connection	Built-in	6.0.0
Micom-VIP	Micom Voice over IP (V/IP)	Built-in	pre-5.0.0
Net2Phone	Net2Phone CommCenter	Built-in	5.2.3
Net2Phone-TCP	Net2Phone Call Setup and Control	Built-in	5.2.3
Net2Phone-UDP	Net2Phone Internet Phone Calls	Built-in	5.2.3
RTCP-B	Real-Time Control Protocol (Broadcast)	Built-in	pre-5.0.0
RTCP-I	Real-Time Control Protocol (Interactive)	Built-in	pre-5.0.0
RTP-B	Real-Time Protocol (Broadcast)	Built-in	pre-5.0.0
RTP-I	Real-Time Protocol (Interactive)	Built-in	pre-5.0.0
SIP	Session Initiation Protocol	Built-in	6.0.0
SIP60	Session Initiation Protocol over port 5060	Built-in	7.1.0
SIP61	Session Initiation Protocol over port 5061	Built-in	7.1.0
SIPX	SIP - Session Initiation Protocol - any port	Built-in	8.2.0
Skinny	Cisco's Skinny Client Control Protocol (SCCP) **Note:** This service is not autodiscovered, but you can create a class manually for this service.	Built-in	6.0.0
Skype	Skype P2P Telephony Application	Built-in	7.0.0
SkypeCommand	Skype Command	Built-in	7.0.0
SkypeData	Skype Data	Built-in	7.0.0
VDOPhone	Service aggregate for Internet telephone (not autodiscovered)	Built-in	pre-5.0.0
VDOPhone-a	Internet telephone application - TCP port 1 (not autodiscovered)	Built-in	pre-5.0.0
VDOPhone-b	Internet telephone application - TCP port 2 (not autodiscovered)	Built-in	pre-5.0.0
VDOPhone-UDP	VDOPhone real-time media (not autodiscovered)	Built-in	pre-5.0.0
Ventrilo	VoIP Group Communications Software	Built-in	8.3.1
Vonage	Vonage VoIP	Built-in	7.1.0
Vonage-SIP	Vonage Session Initiation Protocol	Built-in	7.1.0
Vonage-RTP	Vonage Real-Time Control Protocol	Built-in	7.1.0

URL Categories

The following *category map* (catmap) table lists the categories that PacketShaper can identify and classify in PacketShaper 11.9. The categories are grouped into 17 types, such as adult related, commerce, communication, and so forth. Note that categories periodically change (approximately every 18 months) and your category database will automatically update when this happens. These changes may not be reflected in the documentation, however.

ID	Category Name	Description	Example URLs
Adult-Related Categories			
1	Adult/Mature Content	Sites that contain material of adult nature that does not necessarily contain excessive violence, sexual content, or nudity. These sites include very profane or vulgar content and sites that are not appropriate for children.	humorbomb.com landoverbaptist.org punchbaby.com
93	Sexual Expression	Sites that provide information, promote, or cater to alternative sexual expressions in their myriad forms. Includes but is not limited to the full range of non-traditional sexual practices, interests, orientations or fetishes. This category does not include sites that are sexually gratuitous in nature which would typically fall under the Pornography category, nor does it include lesbian, gay, bi-sexual, transgender or any sites which speak to one's sexual identity.	dailydiapers.com tigerden.com ncsfreedom.org
26	Child Pornography	Sites that include a visual depiction of a minor engaging in sexually explicit conduct.	
7	Extreme	Sites that are extreme in nature and are not suitable for general consumption. Includes sites that revel and glorify in gore, human or animal suffering, scatological or other aberrant behaviors, perversities or debaucheries. It includes visual or written depictions deemed to be of an unusually horrific nature. These are salacious sites bereft of historical context, educational value or artistic merit created solely to debase, dehumanize or shock. Examples would include necrophilia, cannibalism, scat and amputee fetish sites.	
5	Intimate Apparel/Swimsuit	Sites that contain images or offer the sale of swimsuits or intimate apparel or other types of suggestive clothing. Does not include sites selling undergarments as a subsection of other products offered.	victoriassecret.com fredericks.com
6	Nudity	Sites containing nude or seminude depictions of the human body. These depictions are not necessarily sexual in intent or effect, but may include sites containing nude paintings or photo galleries of artistic nature. This category also includes nudist or naturist sites that contain pictures of nude individuals.	bodyscapes.com nudistnews.com

ID	Category Name	Description	Example URLs
50	Mixed Content/Potentially Adult	Sites with generally non-offensive content but that also have potentially objectionable content such as adult or pornographic material that is not organized so that it can be classified separately. Sites that explicitly exclude offensive content are not included in this category.	istockphoto.com images.search.yahoo.com video.google.com youtube.com fotosearch.com imageshack.us metacafe.com imagebam.com break.com
3	Pornography	Sites that contain sexually explicit material for the purpose of arousing a sexual or prurient interest.	playboy.com
4	Sex Education	Sites that provide information (sometimes graphic) on reproduction, sexual development, safe sex practices, sexuality, birth control, tips for better sex, and sexual enhancement products.	scarleteen.com viagra.com
Commerce Categories			
59	Auctions	Sites that support the offering and purchasing of goods between individuals. Does not include classified advertisements.	ebay.com bidfind.com
32	Brokerage/Trading	Sites that provide or advertise trading of securities and management of investment assets (online or offline). Also includes insurance sites, as well as sites that offer financial investment strategies, quotes, and news.	ameritrade.com etrade.com
21	Business/Economy	Sites devoted to business firms, business information, economics, marketing, business management and entrepreneurship. This does not include sites that perform services that are defined in another category (such as Information Technology companies, or companies that sell travel services). Does not include shopping sites.	ge.com sunbeam.com
31	Financial Services	Sites that provide or advertise banking services (online or offline) or other types of financial information, such as loans and insurance. Does not include sites that offer market information, brokerage or trading services.	americafirst.com paypal.com
45	Job Search/Careers	Sites that provide assistance in finding employment, and tools for locating prospective employers.	hotjobs.yahoo.com flipdog.com monster.com
60	Real Estate	Sites that provide information on renting, buying, or selling real estate or properties. Also includes vacation property rentals such as time-shares and vacation condos.	century21.com realtor.com
58	Shopping	Sites that provide or advertise the means to obtain goods or services. Includes sites that list prices or have a clear way to order merchandise and also includes retail shopping outlets, classified advertisements, price comparisons, hot deals, bargains, etc. Does not include sites that can be classified in other categories (such as Vehicles or Weapons).	amazon.com tigerdirect.com

ID	Category Name	Description	Example URLs
66	Travel	Sites that promote or provide opportunity for travel planning, including finding and making travel reservations, sharing of travel experiences (pro or con) vehicle rentals, descriptions of travel destinations, or promotions for hotels/casinos or other travel related accommodations. Mass transit information including but not limited to posting of schedules/fares or any other public transportation-related data are also included in this category.	travelocity.com wynnlasvegas.com orbitz.com hertz.com
67	Vehicles	Sites that provide information on or promote vehicles, boats, or aircraft, including sites that support online purchase of vehicles or parts.	toyota.com autotrader.com boattrader.com
88	Web Ads/Analytics	Sites that provide online advertisements or banners. Does not include advertising servers which serve adult-oriented advertisements. Includes web API embeds (e.g. facebook “like” button).	adserver.yahoo.com adserver.inetzone.com banner.euroads.no
Communication Categories			
51	Chat (IM)/SMS	Sites that provide chat, text messaging (SMS) or instant messaging capabilities or client downloads.	snapchat.com hangouts.google.com chat.com
52	Email	Sites offering web-based email services,such as online email reading, and mailing list services.	mail.yahoo.com outlook.com
53	Newsgroups/Forums	Sites that primarily offer access to newsgroups, messaging or bulletin board systems, or group blogs where participants can post comments, hold discussions, or seek opinions or expertise on a variety of topics.	groups.google.com blog.netscape.com forums.asp.net
110	Internet Telephony	Sites that facilitate Internet telephony or provide Internet telephony services such as voice over IP (VOIP).	speakfreely.org linphone.org skype.com net2phone.com pccall.com twinklephone.com avavoip.com cliconnect.com tivi.com qutecom.org
111	Online Meetings	Sites that facilitate online meetings or provide online meeting, conferencing or training services.	gotomeeting.com beamyourscreen.com dimdim.com webex.com wiredred.com ivisit.com fuzemeeting.com ezmeeting.com gatherplace.net gomeetnow.com
File Transfer Categories			

ID	Category Name	Description	Example URLs
56	File Storage/Sharing	Sites that provide a secure, encrypted, off-site backup and restoration of personal data. These online repositories are typically used to store, organize and share videos, music, movies, photos, documents and other electronically formatted information. Sites that fit this criteria essentially act as your personal hard drive on the Internet.	box.net ibackup.com flickr.com
83	Peer-to-Peer (P2P)	Sites that distribute software to facilitate the direct exchange of files between users. P2P includes software that enables file search and sharing across a network without dependence on a central server.	datorrents.com bittorrentbooster.com frostwire.com teamtptb.com fatal-torrent.com crazy-torrent.com torrenthub.org yourtorrents.info gotorrent.com
71	Software Downloads	Sites that are dedicated to the electronic download of software for any type of computer or mobile device, whether for payment or at no charge.	download.com tucows.com
Health Related Categories			
16	Abortion	Sites which provide information or arguments in favor of or against abortion, describe abortion procedures, offer help in obtaining or avoiding abortion, or provide information on the effects, or lack thereof, of abortion.	gynpages.com abortionfacts.com
23	Alcohol	Sites that offer for sale, promote, glorify, review or in any way advocate the use or creation of alcoholic beverages, including but not limited to beer, wine, and hard liquors. It does not include sites that sell alcohol as a subset of other products such as restaurants or grocery stores.	budweiser.com jackdaniels.com coors.com
37	Health	Sites that provide advice and information on general health such as fitness and well-being, personal health or medical services, prescription medications, alternative and complementary therapies, medical information about ailments, dentistry, optometry, general psychiatry, self-help, and support organizations dedicated to a disease or condition.	cvs.com webmd.com
64	Restaurants/Dining/Food	Sites that list, review, discuss, advertise and promote food, catering, dining services, cooking and recipes.	foodnetwork.com zagat.com
24	Tobacco	Sites that offer for sale, promote, glorify, review or in any way advocate the use or creation of tobacco or tobacco related products including but not limited to cigarettes, pipes, cigars and chewing tobacco. It does not include sites that sell tobacco as a subset of other products such as grocery stores.	bat.com esmokes.com philipmorrisusa.com
Information Related Categories			
27	Education	Sites that offer education information, distance learning, or trade school information or programs. Includes sites that are sponsored by schools, educational facilities, faculty, or alumni groups.	education-world.com ed.gov nyu.edu

ID	Category Name	Description	Example URLs
107	Informational	Sites that provide content that is informational in nature and does not provide a way to directly act upon the information (for example., a site that provides lottery results, but does not sell lottery tickets). This category is always used in conjunction with another category appropriate to the subject matter (for example, gambling).	lotteryusa.com concierge.com wizards.com
46	News/Media	Sites that primarily report information or comments on current events or contemporary issues of the day. Also includes news radio stations and news magazines. Does not include sites that can be rated in other categories.	cnn.com foxnews.com msnbc.msn.com
49	Reference	Sites containing personal, professional, or educational reference, including online dictionaries, maps, censuses, almanacs, library catalogues, genealogy-related sites and scientific information.	dictionary.reference.com encarta.msn.com familysearch.org
40	Search Engines/Portals	Sites that support searching the Internet, indices, and directories.	google.com yahoo.com
95	Translation	Sites that allow translation of text (words, phrases, web pages, between various languages) or that can be used to identify a language.	microsofttranslator.com translate.google.com babelfish.yahoo.com translation.babylon.com worldlingo.com systran.co.uk reverso.net freetranslation.com online-translator.com freedict.com
Leisure Categories			
30	Art/Culture	Sites that nurture and promote cultural understanding of fine art including but not limited to sculpture, paintings and other visual art forms, literature, music, dance, ballet, and performance art and the venues or foundations that support, foster or house them such as museums, galleries, symphonies and the like. Sites that provide a learning environment or cultural awareness outside of the strictures of formalized education such as museums and planetariums are included under this heading.	nga.gov haydenplanetarium.org sydneyoperahouse.com nationalgallery.org.uk poetry.com philorch.org sfballet.org metmuseum.org broadway.com si.edu
20	Entertainment	Sites that provide information on or promote mass entertainment media including but not limited to film, film trailers, television, home entertainment, music, comics, entertainment-oriented periodicals, reviews, interviews, fan clubs, and celebrity gossip.	tmz.com imdb.com marvel.com mtv.com rottentomatoes.com accesshollywood.com celebrity-gossip.net purethenightclub.com greenday.com eonline.com

ID	Category Name	Description	Example URLs
87	For Kids	Sites designed specifically for children. This category is used in conjunction with other categories - it is not a stand-alone category.	kids.yahoo.com playhousedisney.com
33	Games	Sites that support playing or downloading video games, computer games, or electronic games. Includes sites with information, tips or advice on such games or how to obtain cheat codes. Also includes magazines dedicated to computerized games and sites that support or host online sweepstakes and giveaways.	nintendo.com gamespot.com gamedaily.com
68	Humor/Jokes	Sites that primarily focus on comedy, jokes, fun, etc. May include sites containing jokes of adult or mature nature. Sites containing humorous Adult/Mature content also have an Adult/Mature category rating.	ahajokes.com comedycentral.com the-jokes.com
65	Sports/Recreation	Sites that promote or provide information about spectator sports or recreational activities. It does not include sites dedicated to hobbies such as gardening, collecting, board games, scrapbooking, quilting, etc.	espn.com nba.com nfl.com
Liability Concern Categories			
11	Gambling	Sites where a user can place a bet or participate in a betting pool, participate in a lottery, or receive information, assistance, recommendations, or training in such activities. Does not include sites that sell gambling-related products/machines or sites for offline casinos and hotels, unless they meet one of the above requirements.	tipsport.pl downloadcasino.es loteriaypremios.com spil2vind.dk casino-experts.com aceshighcasinoparties.com socialholdemsecrets.com slotmachines.top-webcasinos.com freeslotsworld.net
25	Controlled Substances	Sites that promote, offer, sell, supply, encourage or otherwise advocate the illegal use, cultivation, manufacture, or distribution of drugs, pharmaceuticals, intoxicating plants or chemicals and their related paraphernalia.	sensiseed.com cogollo.net napnt.org magic-mushroom.com
121	Marijuana	Sites that discuss, encourage, promote, offer, sell, supply or otherwise advocate the use, cultivation, manufacture or distribution of marijuana and its myriad aliases, whether for recreational or medicinal purposes. Includes sites with content regarding marijuana-related paraphernalia.	hightimes.com rollitup.org grasscity.com cannabis.com thctalk.com cannabisculture.com
9	Scam/Questionable/Illegal	Sites that promote scams or that provide or sell legally questionable educational materials such as term papers. Also includes sites that advocate or give advice on performing acts that are illegal or of questionable legality such as service theft, evading law enforcement, fraud, burglary techniques and plagiarism.	oppapers.com antiessays.com gregmiller.net/locks/mitguide/ cleaneasy.com

ID	Category Name	Description	Example URLs
14	Violence/Hate/Racism	Sites that depict extreme physical harm to people, animals or property, or that advocate or provide instructions on how to cause such harm. Also includes sites that advocate, depict hostility or aggression toward, or denigrate an individual or group on the basis of race, religion, gender, nationality, ethnic origin, or other involuntary characteristics. Includes content that glorifies self-mutilation or suicide.	armyofgod.com gorillafights.com torturereom.com nutabuse.com anti-bush.com free-gaza.com teengirlfights.net klanparenthood.com newnation.org festered.com
15	Weapons	Sites that sell, review, or describe weapons such as guns, knives or martial arts devices, or provide information on their use, accessories, or other modifications. Does not include information on BB guns, paintball guns, black powder rifles, target shooting, or bows and arrows unless the site also meets one of the above requirements. Also does not include sites that promote collecting weapons, or groups that either support or oppose weapons use.	browning.com remington.com
Multimedia Categories			
84	Audio/Video Clips	Sites that provide streams or downloads of audio or video clips-typically less than 15 minutes or less in length. Also includes sites that provide downloaders and players for audio and video clips.	new.music.yahoo.com music.msn.com vodpod.com funnyordie.com video.nationalgeographic.com videocure.com muchmusic.com soundboard.com video-clips.co.uk cmt.com
112	Media Sharing	Sites that allow sharing of media (e.g., photo sharing) and have a low risk of including objectionable content such as adult or pornographic material.	shutterfly.com zoomonga.com fotki.com www.pixagogo.com kodakgallery.com photos.google.com snapfish.com dropshots.com smugmug.com webshots.com

ID	Category Name	Description	Example URLs
113	Radio/Audio Streams	Sites that provide streams or downloads of radio, music, or other audio content-typically more than 15 minutes in length.	pandora.com shoutcast.com last.fm napster.com grooveshark.com deezer.com slacker.com radiopaq.com knac.com jazz24.org
114	TV/Video Streams	Sites that provide streams or downloads of television, movie, Webcam, or other video content-typically more than 15 minutes in length.	hulu.com fancast.com tvduck.com casttv.com demand.five.tv tv.com netflix.com sling.com thewb.com vudu.com
Security Concern Categories			
17	Hacking	Sites that distribute, promote or provide hacking tools and/or information which may help gain unauthorized access to computer systems and/or computerized communication systems. Hacking encompasses instructions on illegal or questionable tactics, such as creating viruses, distributing cracked or pirated software or other protected intellectual property.	astalavista.box.sk happyhacker.org flyninja.net 1nj3ct.in
118	Piracy/Copyright Concerns	Sites that provide information or technology for cracking or pirating software or other protected intellectual property, and sites that distribute such media.	megashare.info fullepisode.info thepiratebay.se isohunt.com gamecopyworld.com cracks4u.us eztv.it watchseries.lt
108	Computer/Information Security	Sites that provide information or tools for securing or safeguarding computers, networks, and other data systems. While these sites provide helpful and legitimate security information to IT professionals, they also pose a degree of risk because information they provide may be used to help gain unauthorized access to systems.	metasploit.com cve.mitre.org nessus.org sectools.org trustedsec.com itsecurity.com
98	Placeholders	Sites that are under construction, parked domains, search-bait or otherwise generally having no useful value.	
102	Potentially Unwanted Software	Sites that distribute software that is not malicious but may be unwanted within an organization such as intrusive adware and hoaxes.	

ID	Category Name	Description	Example URLs
57	Remote Access Tools	Sites that primarily focus on providing information about and/or methods that enable authorized access to and use of a desktop computer or private network remotely.	gotomypc.com
101	Spam	Sites that are part of the spam ecosystem, including sites linked in unsolicited bulk electronic messages and sites used to generate or propagate such messages.	
92	Suspicious	Sites considered to have suspicious content and/or intent that poses an elevated security or privacy risk. This categorization is determined by analysis of web reputation factors. If a site is determined to be clearly malicious or benign, it will be placed in a different category.	
Security Threat Categories			
18	Phishing	Sites that are designed to appear as a legitimate bank or retailer with the intent to fraudulently capture sensitive data (i.e. credit card numbers, pin numbers).	
86	Proxy Avoidance	Sites that provide information on how to bypass proxy server features or gain access to URLs in any way that bypasses the proxy server. This category includes any service which will allow a person to bypass the Symantec filtering system, such as anonymous surfing services.	anonymizer.com
44	Malicious Outbound Traffic / Botnets	Sites to which botnets or other malware (as defined in the Malicious Sources category) send data or from which they receive command-and-control instructions. Includes sites that contain serious privacy issues, such as "phone home" sites to which software can connect and send user information. Usually does not include sites that can be categorized as Malicious Sources.	
43	Malicious Sources/Malnets	Sites that host or distribute malware or whose purpose for existence is as part of the malware ecosystem. Malware is defined as software that takes control of a computer, modifies computer settings, or collects or reports personal information without the permission of the end user. It also includes software that misrepresents itself by tricking users to download or install it or to enter personal information. This includes sites or software that perform drive-by downloads; browser hijackers; dialers; any program that modifies your browser homepage, bookmarks, or security settings; and keyloggers. It also includes any software that bundles malware (as defined above) as part of its offering. Information collected or reported is "personal" if it contains uniquely identifying data, such as email addresses, name, social security number, IP address, etc. A site is not classified as malware if the user is reasonably notified that the software will perform these actions (e.g., it alerts that it will send personal information, be installed, or that it will log keystrokes).	screensaver.com
Social Interaction Categories			

ID	Category Name	Description	Example URLs
63	Personal Sites	Sites that primarily offer access to personal pages and blogs. This classification includes but is not limited to content that shares a common domain such as Web space made available by an ISP or some other hosting service. Personal home pages and blogs tend to be dynamic in nature and their content may vary from innocuous to extreme.	blogger.com wordpress.org
106	E-Card/Invitations	Sites that facilitate the sending of electronic greeting cards, animated cards, or similar electronic messages typically used to mark an event or occasion.	123greetings.com bluemountain.com greetingcards.com americangreetings.com egreetings.com memorygreetings.com ecardfunny.com e-cards.com pencake.com regards.com
47	Personals/Dating	Sites that promote interpersonal relationships.	eharmony.com lovesites.com dating.com
55	Social Networking	Sites that enable people to connect with others to form an online community. Typically members describe themselves in personal Web page profiles and form interactive networks, linking them with other members based on common interests or acquaintances. Instant messaging, file sharing and Web logs (blogs) are common features of Social Networking sites. These sites may contain offensive material in the community-created content. This category may be used in conjunction with another category for more narrowly-focused social networking sites, such as professional networking sites or social networking sections of Personals/Dating sites. **Note:** The Social Networking category includes classification of sites that may not yet have a service in PacketShaper. Therefore, creating a Social_Networking category class may be preferable over creating a SocialNetwork service group class.	facebook.com instagram.com plus.google.com

Society/Government Categories

ID	Category Name	Description	Example URLs
22	Alternative Spirituality/Belief	Sites that promote and provide information on a wide range of non-traditional and/or non-religious spiritual, existential, experiential, and philosophical belief systems. Includes sites related to atheism, agnosticism, and mysticism; sites related to quasi-religious, philosophical or spiritual belief systems and practices that do not include formally established religious meetings, places of worship, organizational structure, or dogma; and sites that endorse or offer information about affecting or influencing real events through supernatural or magical means. Also includes sites that discuss or deal with paranormal or unexplained events. This category does not include sites centered around traditional, organized religious belief, practice, and observance (Religion).	atheists.org ufocasebook.com paranormalnetwork.net ancientblackmagic.com spellsandmagic.com churchofsatan.com paranormalnews.com morespells.com nuforc.org neworleansghosts.com
29	Charitable Organizations	Sites that foster volunteerism for charitable causes. Also encompasses non-profit associations that cultivate philanthropic or relief efforts. Does not include organizations that attempt to influence legislation as a significant portion of their activities or organizations that campaign for, contribute to or affiliate with political organizations or candidates.	scouting.org 4-h.org ymca.net lionsclubs.org toastmasters.org redcross.org unicef.org pewtrusts.org childrensaidsociety.org cityharvest.org
34	Government/Legal	Sites sponsored by or which provide information on government, government agencies and government services such as taxation and emergency services. Includes sites which discuss or explain laws of various governmental entities. Also includes sites that advertise legal services, lawyers for hire, adoption services, information about adoption, immigration information, and immigration services.	whitehouse.gov federalreserve.gov
35	Military	Sites that promote or provide information on military branches or armed services.	www.army.mil www.navy.mil www.af.mil
36	Political/Social Advocacy	Sites sponsored by or which provide information on political parties, special interest groups, or any organization that promotes change or reform in public policy, public opinion, social practice, or economic activities.	texasgop.org aclu.org rnc.org dnc.org
54	Religion	Sites that promote and provide information on traditional, organized religious belief, practice and observance and directly-related subjects such as religious catechism or dogma and places of religious worship or observance (e.g., churches, synagogues, temples, etc.). This category does not include sites about non-traditional spiritual and non-religious belief systems (Alternative Spirituality/Belief).	catholic.net gospel.com lds.org

ID	Category Name	Description	Example URLs
61	Society/Daily Living	Sites providing information on matters of daily life. This includes but is not limited to pet care, home improvement, fashion/beauty tips, hobbies and other tasks that comprise everyday life. It does not include sites relating to entertainment, sports, jobs, personal pages or other topics which already have a specific category.	lifestyle.msn.com hobbies.net akc.org
Technology Categories			
38	Technology/Internet	Sites that sponsor or provide information on computers, technology, the Internet and technology-related organizations and companies.	dell.com microsoft.com javaworld.com
97	Content Servers	Servers that provide commercial hosting for a variety of content such as images and media files. These servers are typically used in conjunction with other web servers to optimize content retrieval speeds.	akamai.com cdnetworks.com limelightnetworks.com
103	Dynamic DNS Host	Sites that utilize dynamic DNS services to map their domain names to dynamic IP addresses.	dyndns.com no-ip.com
96	Non-viewable/Infrastructure	Servers with non-malicious, non-offensive content or resources used by applications, but not directly viewable by web browsers. Includes but is not limited to Web analytics sites (such as visitor tracking and ranking sites) and content filtering systems. Also includes third-party authentication (Google/Twitter/Facebook auth).	
85	Office/Business Applications	Sites with interactive, Web-based office/business applications. This excludes email, chat/IM or other sites that have a specific content category.	
89	Web Hosting	Sites of organizations which provide top-level domain pages, as well as web communities or hosting services. Includes blog hosting sites.	sites.google.com 1and1.com
109	Internet Connected Devices	Sites that allow management and monitoring of or network access to physical devices connected to the Internet. Such devices include but are not limited to network infrastructure such as routers and switches, network-enabled industrial equipment, security cameras, home automation equipment, and other Web-enabled devices. Also includes security camera feeds, which are dually categorized as TV/Video Streams.	online.wilife.com mydlink.com arc.webeyecms.com viewnetcam.com synology.me mycloudnas.com customer.control4.com webcam.oii.ox.ac.uk

Preconfigured Graphs (Advanced UI)

This section describes all the graphs available in the *Statistics:Reports* window. Click a graph name below to see a description and a sample graph.

These graphs are grouped into the following basic categories, or types of graphs:

- Bandwidth Usage
- Efficiency
- Top-10
- Response Time Measurement
- TCP Connections

Graphs for Analyzing Bandwidth Usage

Graph Name	When to Use it
Class Utilization	To analyze a class' average bandwidth usage
Class Utilization with Peaks	To see if a class is frequently hitting a capacity limit
Dynamic Partition Usage	To determine whether the size of a dynamic partition is appropriate
Link Utilization	To determine whether the link usage varies a lot and to see what your average capacity needs are To see whether the link size has changed over the interval (if you choose to display the link size when creating the graph) To examine the utilization of an inbound or outbound link
Link Utilization with Peaks	To compare a link's average and peak bandwidth usage with the link size so that you can determine whether the link size is sufficient (if you choose to display the link size when creating the graph)
Partition Utilization	To examine a partition's average bandwidth usage To compare a partition's average bandwidth usage with the partition's size so that you can determine whether the size is appropriate (if you choose to display the partition size when creating the graph)
Partition Utilization with Peaks	To determine the average and peak bandwidth usage of a partition To compare a partition's average and peak bandwidth usage with the partition's size (if you choose to display the partition size when creating the graph)

Graphs to Analyze Efficiency

Graph Name	When to Use it
Bytes Transmitted	To determine the actual volume (such as number of gigabytes) that has been transmitted in a time period for a specific link, partition, or class
Guaranteed Rate Failures	To see the number of times PacketShaper was unable to provide the bandwidth guaranteed by the class' policy; allows you to analyze the effectiveness of a policy
Network Efficiency	To see what percentage of your network traffic is devoted to packet retransmissions. Actually, this graph shows the percentage of packets that are *not* retransmissions, so make sure you interpret this graph correctly. For example, if network efficiency is 90%, then 10% of the packets are retransmissions.
Packet Size Distribution	To analyze the distribution of packet sizes on an Inbound or Outbound link
Packets Transmitted	To determine the cause of a network slowdown. Compare the Packets Transmitted chart to the corresponding class or link utilization chart to determine if a link or class is being flooded by small packets. The flooding of small packets reduces the rate at which full-size packets get transmitted, which could slow down the network.

Graphs for Analyzing Top Ten

Graph Name	When to Use it
Top 10 Partitions	To quickly see which partitions are consuming the most bandwidth on the network
Top 10 Classes	To quickly see which classes are consuming the most bandwidth on the network
Top 10 Children Classes	To see the biggest consumers of bandwidth for a class' direct descendant children

Graphs for Analyzing Response Time

Graph Name	When to Use it
Average Transaction Size	To view information about transaction size as part of interpreting RTM data. For example, you can compare a Network Delay graph and an Average Transaction Size graph for the same time period to get a better picture of what is causing a delay.
Network Delay	To analyze the portion of transaction time that is attributable to the network. Since network delay is affected by transaction size, the Normalized Network Delay graph could be more useful when there is high variability in transaction size.
Network Delay Distribution	To see the breakdown of network delays for a class, partition, or link

Graph Name	When to Use it
Normalized Network Delay	To analyze response-time data accounting for variation in transaction sizes.
Packet Round-Trip Time	To track the history of network speed, independent of the size of the transactions
Server Delay	To analyze the portion of transaction time that is attributable to the server
Server Delay Distribution	To see the breakdown of server delays for a class, partition, or link
Service Level Compliance	When you've set a total delay threshold and service level threshold, you can see how a class is performing with respect to the thresholds; the graph shows the percentage of transactions whose delays are in an acceptable range
Transaction Delay	To determine the source (server or network) of any significant delays that a traffic class experiences; to compare actual response times to the total delay threshold
Transaction Delay Distribution	To get a more complete picture of transaction delays; this graph shows the number of transactions whose total delay falls into various transaction-time buckets; for example, you can see that most transactions experience .1 ms or .25 ms total delay and only a few experience 1 ms.

Graphs for Analyzing TCP Connections

Graph Name	When to Use it
Connection Retransmissions	To identify poorly tuned servers whose retransmissions are wasting bandwidth; to show applications that could benefit from traffic shaping or rate control
TCP Connections Initiated	To view the rate at which new connections are opened
TCP Health	To view a comprehensive picture of TCP connections for a class, partition, or link. Use this graph to identify abnormal network conditions, analyze why a server isn't performing properly, or determine whether a DoS attack occurred.

Sky UI Graphs

Sky UI offers the following types of graphs for displaying historical class data. See Measurement Variables for descriptions of the variables used in the graphs.

Graph Name	When to Use it	Measurement Variables
Class Utilization, Average	To analyze the average bandwidth usage of the link, a class, or several classes	avg-bps
Class Utilization, Peaks	To analyze peak bandwidth usage of the link, a class, or several classes to see if a capacity limit is frequently being hit	peak-bps
Guaranteed Rate Failures	To see the number of times PacketShaper was unable to provide the bandwidth guaranteed by the class' policy; allows you to analyze the effectiveness of a policy	guar-rate-fails
Network Efficiency	To see what percentage of your network traffic is devoted to packet retransmissions. Actually, this graph shows the percentage of packets that are *not* retransmissions, so make sure you interpret this graph correctly. For example, if network efficiency is 90%, then 10% of the packets are retransmissions.	tcp-efficiency%
Bytes Transmitted	To determine the actual volume (such as number of gigabytes) that has been transmitted in a time period for a specific class	bytes tcp-retx-bytes
Packets Transmitted	To determine the cause of a network slowdown. Compare the Packets Transmitted chart to the corresponding class utilization chart to determine if a class is being flooded by small packets. The flooding of small packets reduces the rate at which full-size packets get transmitted, which could slow down the network.	pkts tcp-retx-pkts
TCP Health	To view a comprehensive picture of TCP connections for a class. Use this graph to identify abnormal network conditions, analyze why a server isn't performing properly, or determine whether a DoS attack occurred.	tcp-conn-inits tcp-conn-aborts tcp-conn-server-ignores tcp-conn-server-refuses
TCP Connections Initiated	To view the rate at which new connections are opened	tcp-conn-inits
Connection Retransmissions	To identify poorly tuned servers whose retransmissions are wasting bandwidth; to show applications that could benefit from traffic shaping or rate control	tcp-retx-pkts tcp-early-retx-toss-pkts%
Transaction Delay	To determine the source (server or network) of any significant delays that a traffic class experiences; to compare actual response times to the total delay threshold	total-delay-avg network-delay-avg server-delay-avg total-delay-threshold

Graph Name	When to Use it	Measurement Variables
Normalized Network Delay	To analyze response-time data accounting for variation in transaction sizes.	normalized-network-delay-avg
Packet Round Trip Time	To track the history of network speed, independent of the size of the transactions	avg-round-trip-time

PacketShaper Port and URL Usage

PacketShaper uses the following ports while operating. Ensure that your firewall allows these ports.

Inbound Connections to PacketShaper

Whether for administrative access, or to accept incoming data to be managed, this table details the connection points that are open on PacketShaper.

Service	Port	Protocol	Configurable?	Source	Description
HTTP	80	TCP	no	user's client	Web service for PacketShaper Sky and Advanced UI
HTTPS	443	TCP	no	user's client	Secure web service for PacketShaper Sky and Advanced UI
NTP	123	UDP	yes	time server	Synchronize with time servers
Secure Shell (SSH)	22	TCP	no	user's client	Securely manage and configure PacketShaper with a command line interface.
SNMP	161	UDP	no	SNMP analysis tools	Listen for queries from remote SNMP analysis tools (if SNMP is enabled).
Standby	2014	TCP	no	standby partner	Standby partner communication

Outbound Connections from PacketShaper

PacketShaper connects to the services listed below. Note that many of these services are optional, and the ports don't need to be open on the firewall unless they are being used.

Service	Port	Protocol	Configurable?	Destination	Function
BCAAA	16101	TCP	yes	BCAAA server on Active Directory	Look up user names and groups on Symantec Authentication and Authorization Agent server.
DNS	53	TCP/UDP	no	DNS server	Perform domain name resolution for URLs in data sent to PacketShaper for scanning, and to resolve Internet addresses the appliance connects to.
FDR	9800	UDP	yes	FDR collector	Send flow detail records to FDR collector
Web Proxy	*user defined*		yes	Web proxy server	All PacketShaper features that access external servers on the Internet will go through the proxy server. This server handles WebPulse requests, category map downloads, heartbeat emissions, support status updates, and image updates.
PolicyCenter	*user defined*	TCP	yes	PolicyCenter appliance	Share configuration with PolicyCenter appliance.
RADIUS Authentication	1812	TCP/UDP	yes	RADIUS authentication server	Communicate with RADIUS servers to authenticate PacketShaper administrators

Service	Port	Protocol	Configurable?	Destination	Function
RADIUS Accounting	1813	TCP/UDP	yes	RADIUS accounting server	Communicate with RADIUS accounting servers to have an audit trail for user logins.
SMTP	25	TCP	yes	Mail server	Send email notifications.
SNMP	162	UDP	yes (SNMPv3)	Trap receiver	Send SNMP traps.
Syslog	514	UDP	yes	Syslog server	Report appliance health and statistical data to a syslog server.
TACACS	49	TCP/UDP	yes	TACACS+ server	Communicate with TACACS+ servers to authenticate PacketShaper administrators and/or produce an audit trail for user logins.

Required URLs

Ensure connectivity to the following URLs, used by PacketShaper features.

URL	Protocol	Port	Description
bto.bluecoat.com	https/TCP	443	Support links to software, support cases and documentations
subscription.es.bluecoat.com	https/TCP	443	Symantec licensing
sp.cwfservice.net	https/TCP	443	WebPulse update server
sitereview.bluecoat.com	https/TCP	443	WebPulse map update server
hb.bluecoat.com	https/TCP	443	Symantec heartbeat server
cda.bluecoat.com	https/TCP	443	Traffic information reporting server
updates.bluecoat.com	https/TCP	443	Support update server
time.nist.gov*	UDP	123	NTP server (primary)
time-a.nist.gov*	UDP	123	NTP server (secondary)

Note: * NTP servers and ports are user-configurable; default values are shown here.

Glossary of Terms

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

9.1026/	A top-level folder on the PacketShaper that contains files that support the URL category and web application features.
9.256/	The top-level folder that contains configuration, log, and command files.
9.258/	The top-level folder that contains packet logs, images, diagnostic files, and measurement data.
Access Link	The WAN connection that the PacketShaper manages. See also Link and Link Rate.
Action File	Command files that are associated with a specific adaptive response agent. See Action Files Overview.
Adaptive Response	A PacketShaper feature that monitors unit, application, and network health, providing a color-coded summary that lets you identify potential problems. See Adaptive Response Overview.
Address Resolution Protocol (ARP)	A TCP/IP protocol used to map an IP address to a physical hardware address.
Admission Control	A mechanism for a rate policy that determines what happens when there isn't enough bandwidth to satisfy guaranteed rate requests (such as refuse connections or give a trickle of bandwidth). See Apply Admission Control.
Agent	An entity in adaptive response that measures and monitors performance of the PacketShaper, the network, or a specific class. See Adaptive Response Overview.
ARP	See Address Resolution Protocol.
Asymmetric Traffic	Flows that enter the network via one path and exit via a different path.
BCAAA	Blue Coat Authentication and Authorization Agent (BCAAA) converts IP addresses to user names so that it can classify and report on network users. See BCAAA Overview.
Burstable Partition	A type of partition that allows an aggregate traffic class to use a defined amount of bandwidth, and also allows that class to access additional unused bandwidth, if needed.
Bypass Mode	Operational state in which the appliance connects the Inside and Outside Ethernet ports directly to each other, essentially behaving like a piece of cable. By default, the unit enters bypass mode on loss of power.
Category	See URL Category.
CIDR	See Classless InterDomain Routing.
Class	A logical grouping of traffic flows that share the same characteristics a specific application, protocol, address, or set of addresses. Also known as *traffic class*.
Classless InterDomain Routing (CIDR)	A scheme for dividing the network and host portions of an IP v4 address on a bit-by-bit basis. Replaces the original byte-delimited class A,B,C,D IP address method, and provides greater flexibility in dividing a finite number of IP addresses among many subnets.
CLI	See **Command-Line Interface**.
Cluster	A group of PacketShapers that have been configured for high availability.
Collector, Flow Detail Record	A software application that accumulates the data from an FDRemitter (PacketShaper). Most collectors do much more than gather the data they also present the information in a meaningful way in reports and graphs.
Command-Line Interface (CLI)	The terminal access method, available via remote login or direct connection of an ASCII terminal to the Console port. Commands are entered as text on a single line.

Concurrent Hosts	IP hosts that the PacketShaper has "learned." For the maximum number of concurrent hosts that you can have on a unit, see Configuration Limits for PacketShaper. See also Host Database and sys limits.
Diffserv	Differentiated Services a protocol for specifying and controlling network traffic so that certain types of traffic get precedence. See also DSCP.
Discard Policy	A bandwidth allocation rule that blocks all packets for a traffic type.
Domain	A unique identifier for a group of network resources, often representing a specific organization for example, *symantec.com*.
Domain Name Server (DNS)	A network computer that translates a text-based name (such as www.symantec.com) for a network resource to the dotted-decimal IP address (such as 10.10.10.10) required to access the resource.
Domain Name	A text-based name for an Internet address for example, *www.symantec.com*. A client computer contacts a Domain Name Server to request translation of the human-readable domain name into the dotted-decimal address format used by the Internet Protocol.
DRTR	Dynamic Real Time Rating service. A service that determines the category of an unrated website by retrieving key pieces of the page's content and context.
DSCP	Differentiated Services Code Point—A six-bit field in the IP header that specifies the per hop behavior for a given flow of packets. See Classify Diffserv Code Point Traffic and Substitute Diffserv Values.
Dynamic Partition	A type of partition that automatically creates subpartitions on the fly as users become active in a traffic class. This capability allows service providers or enterprise customers to guarantee a user a minimum amount of bandwidth at all times.
Emitter, Flow Detail Record	A device that pushes data to a remote system (collector) on a continuous basis. PacketShaper is an emitter of flow detail records.
Ethernet	A protocol for electrical connections and base addressing that is widely used in local area networks.
Event	In PacketShaper, a defined occurrence, such as an application reaching a certain bit rate, that can be used to trigger an alert message sent via email, an SNMP trap message, or both.
Exception Class	A class that has been tagged to be sorted above standard (non-exception) classes in the traffic tree. Exception classes give you the ability to redefine the search order that PacketShaper uses to find a match for a traffic flow.
Exception Host List	A host list that contains a list of IP addresses you don't want an adaptive response agent to monitor (for example, servers). See Agent Host Lists.
Excess Rate	The portion of the data rate for a traffic class that exceeds the minimum allocated rate.
Flash Memory	A type of nonvolatile, electronic storage media used in the PacketShaper to store the PacketShaper software. Also called system memory.
Flow	A specific instance of a connection, session, or packet-exchange activity by an identified traffic class.
Flow Detail Record (FDR)	Information about a TCP or non-TCP flow, such as where the flow originated and where it went to, the size of the flow (in terms of packets and bytes), and when the flow was sent. For more information, see Flow Detail Records Overview.
Full Duplex	A mode of data transmission in which each device is able to transmit and receive data simultaneously. See also Half Duplex.
Gateway	The network device, often a router, to which packets destined for another network are sent.
Guaranteed Rate	In a rate policy, the minimum allocated bandwidth made available to each connection in a traffic class. Unused bandwidth is loaned to other traffic classes.
Guaranteed Rate Failure	An instance when PacketShaper is unable to allocate the guaranteed rate when applying a class' rate policy.
Half Duplex	A mode of data transmission in which each device alternately receives and transmits data.

Heartbeat	A message periodically emitted from the PacketShaper to a designated Symantec server; this message provides Symantec with useful support-related information.
High Availability	A network topology feature that ensures mission critical applications are available 100% of the time, on a 24x7 basis. This goal is typically accomplished by having multiple access routers with multiple WAN link interfaces. PacketShapers can enhance a high availability topology by taking over traffic management on a disabled link.
Hits	The number of times traffic flows match the criteria specified in a traffic class.
Host	A computer that is connected to a TCP/IP network, including the Internet. Each host has a unique IP address.
Host Accounting	A feature that allows you to track byte counts of throughput for each IP address, group of users in a host list, or subnet. See Host Accounting.
Host Database (hostdb)	A record of all hosts that have active connections through the PacketShaper. Once a host closes its connection, the host will be purged from the database. In addition, the unit will clear host entries if they aren't active for approximately ten minutes. Thus, the hostdb is a real-time list of hosts. The PacketShaper features that use the hostdb include Top Talkers/Listeners, host accounting, and the classification-acceleration cache.
Host List	A host list contains a set of IP addresses, DNS names, and/or subnets. Host lists are useful when creating classes based on hosts, retrieving host accounting data, and defining exception lists for adaptive response host agents.
HTTP over SSL (HTTPS)	A protocol for transferring private documents over the Internet.
Ignore Policy	A bandwidth allocation rule that exempts a traffic class from bandwidth management. Used for traffic not destined for the WAN.
Inbound Rate	In PacketShaper, the maximum rate at which data is arriving at the local site across the wide area network.
Incident Report	A report (such as a graph) that an adaptive response agent automatically generates when it crosses from a green/yellow threshold to a red threshold or from red/yellow to green. The report shows supporting data for the problem at the time it occurred. See View Incident Reports.
Inheritable Traffic Class	A traffic class with a policy that can be given to a sibling class, if the sibling doesn't have its own policy.
Inside NIC Mode	In PacketShaper, the speed selection for the LAN connection facing the local computer network.
Inside Port	The unit's 8-pin modular Ethernet connector for the local LAN.
Internet Protocol (IP)	A network protocol for addressing and sending data packets. The messenger-oriented component of TCP/IP.
IP Address	The dotted-decimal identifier used by the Internet Protocol to contact a specific network device for example, *10.10.10.10*.
IP Precedence	A QoS traffic priority method for TCP/IP that uses the three-bit Type of Service (TOS) field in the IP header.
IQOS	Stands for *Internet Quality of Service*. Used by Symantec to uniquely name LDAP objects created by PacketShaper.
Leaf Class	A child class that doesn't have any of its own child classes.
Legacy Flows	Traffic using a non-TCP/IP protocol, often encapsulated in a TCP or IP wrapper.
Link	A communications connection that is used to transmit data from a source to a destination (for example, a T1 line). See also Access Link and Link Rate.
Link Rate	The capacity of the access link, measured in bits per second. For example, a T1 access link has a link rate of 1.5 Mbps. The link has an inbound and an outbound rate.
Link State Mirroring	When this feature is enabled, PacketShaper will bring down the second port of a NIC pair if the first goes down.
Load Balancing	The process of dividing network traffic between parallel paths or devices in order to handle more transactions without overloading a specific resource.

Local Area Network	A computer network that spans a relatively small area, such as a single building or group of buildings.
Local Host	In networking, a reference to *this* computer. In the traffic tree, the Localhost class refers to the PacketShaper unit.
Look Access	Read-only permission for the PacketShaper software. See also Touch Access.
Look Community String	The PacketShaper password required before an SNMP manager can read the value of a PacketShaper manageable object.
Look Password	The PacketShaper password required before look access is granted.
MAC Cache Entries	Short-term memory records of the data link control layer addresses for known hosts.
Matching Rule	A set of characteristics that identifies a specific traffic type.
Measurement Engine	A PacketShaper background process that collects time-series and histogram data. This data is stored in the unit's 9.258/me/measure folder.
MTU	Maximum Transfer Unit – the largest datagram than can be transmitted by an IP interface, without it needing to be broken down into smaller units.
Net Mask	The dotted-decimal expression that determines which portion of an IP address represents the network address and which portion is for the host for example, *255.0.0.0*.
Network Delay	The number of milliseconds spent in transit when a client and server exchange data. Includes the transit time for all packets required for a request-response transaction.
Network Efficiency	The ratio of TCP packets not requiring retransmission to the total number of TCP packets sent.
Network Interface Card (NIC)	An onboard chipset or an add-in circuit board that provides support for connection to a local area network.
Never-Admit Policy	A bandwidth allocation rule that tells PacketShaper to enforce admission control at the beginning of each flow. Used to reject or redirect web traffic and restrict non-TCP traffic.
New Flows Per Minute	The number of new flows initiated *from* a host (in the case of a client host) or *to* a host (in the case of a server) during a one-minute period.
NIC	See Network Interface Card.
Outbound Rate	In PacketShaper, the maximum rate at which data can be transmitted to another site across the wide area network.
Outside NIC Mode	In PacketShaper, the speed selection for the LAN connection facing the router or WAN connection.
Outside Port	The unit's 8-pin modular Ethernet connector for the path to the router and WAN.
Oversubscribed	A situation in which the sum of all partition sizes on a unit is greater than the link size.
Packet Capture	A feature that captures packets for future analysis, allowing you to analyze detailed information about the packets in a class, such as the source and destination IP addresses and protocols used.
PacketShaper	A Symantec product that is an application-based traffic and bandwidth management system.
Partition	A bandwidth pipe assigned to a given traffic class to protect or restrict all the flows in that class.
Partners	PacketShapers that are members of a high availability cluster.
Plug-Ins	Downloadable classification modules that you can install into an existing software release. Plug-ins are not supported in PacketShaper 11.x.
Policy	A rule assigned to a given traffic class that defines how a single flow will be handled during bandwidth allocation.

PolicyCenter	A Symantec appliance that centralizes management and propagates configuration changes to large numbers of PacketShapers. See PolicyCenter Overview.
Priority Policy	A bandwidth allocation rule that sets the relative precedence for traffic without allocating a specific rate.
Protocol	A set of rules governing the transmission and receiving of data.
ProxySG	A Symantec appliance that secures web communications and accelerates the delivery of business applications.
Proxy Service	A network traffic relay application that represents trusted, usually local, clients when they access unknown network resources. Often used to protect a local corporate LAN from potentially hostile outside hosts.
QoS	See Quality of Service.
Quality of Service (QoS)	In PacketShaper, a measure of the performance of a network component against expectations.
Queuing	The process of delaying packets until a bandwidth-related condition is met.
Rate Policy	A bandwidth allocation rule used to smooth bursty traffic, such as HTTP, using TCP rate control. Rate policies can have a guaranteed minimum rate and/or be burstable at a selected priority level.
Reboot	The process of interrupting a computing device, causing it to reload its operating system and restart all default operations.
Recovery User	If the PacketShaper is continuously rebooting and does not present the login prompt, you can console connect to the appliance and log in as the recovery user. As a recovery user, you can restore to factory defaults and reconfigure the appliance.
Reset	The process of interrupting a computing device and causing it to return to a default condition.
Response Time Measurement (RTM)	A PacketShaper feature that tracks delay statistics for connection-based TCP traffic classes and breaks each response time measurement into network delay, server delay, and total delay. For example, you can measure response times for applications, web traffic, individual hosts, or subnets. See " RTM Overview " on page 104.
RTM	See Response Time Measurement.
SameSide	A special traffic class that measures traffic not destined for the WAN.
Secure Shell (SSH)	A program and protocol that provides strong authentication and secure communications for logging onto a remote computer.
Server Delay	The time the server uses to process a client's request after it receives all required data; the time between when the server receives the last request packet and when it sends the first packet of the reply.
Service Group	A set of services that have been grouped together based on a common functionality; it offers a quick way to classify and control a group of related services. See Service Group Tasks.
Shaping	A PacketShaper bandwidth management feature that improves upon default bandwidth allocation by enforcing policies on traffic flows and partitions on aggregate flows in a traffic class.
Simple Network Management Protocol (SNMP)	A widely used member of the TCP/IP suite of protocols designed to support remote management of network components and the gathering of alarm and event information.
Simple Network Time Protocol (SNTP)	A connectionless (UDP) protocol used to synchronize the time in a client (PacketShaper) to a high accuracy network time server.
SNMP	See Simple Network Management Protocol.
SNMP Traps	Unsolicited, connectionless messages sent by PacketShaper to a configured SNMP host. These messages report alarms or indicate that a configured event has occurred.

SNTP	See Simple Network Time Protocol.
Software Version	The revision of software currently operating in the PacketShaper unit. This information is located on the **info** tab in the PacketShaper browser interface.
SSH	See Secure Shell.
Standby	PacketShaper's feature for high availability and clustering technology.
Standby Partner	See Partners.
Synthetic Transactions	Web or other TCP transactions that the PacketShaper can send at regular intervals to verify the availability of critical hosts.
System Memory	A type of nonvolatile, electronic storage media used in the PacketShaper to store the PacketShaper software.
System Variables	User-adjustable PacketShaper settings. See Adjust System Variables.
TCP	See Transmission Control Protocol.
TCP Flows	The number of unique sessions using Transmission Control Protocol.
TCP Rate Control	A method of throttling traffic flow, specified in the Transmission Control Protocol, and utilized by PacketShaper.
Template, Adaptive Response	A form that is used as a guide for creating adaptive response agents. You create an agent by selecting a template and defining the agent parameters. For a list of templates, see Agents Overview.
Top Listeners	An identified traffic class for which PacketShaper has been configured to record the host names or IP addresses of the devices receiving the greatest amounts of traffic (the "listeners").
Top Talkers	An identified traffic class for which PacketShaper has been configured to record the host names or IP addresses of the devices transmitting the greatest amounts of traffic (the "talkers").
Touch Access	Read/write permission for the PacketShaper software. See also Look Access.
Touch Community String	The PacketShaper password required before an SNMP manager can set the value of a PacketShaper manageable object.
Touch Password	The password required by PacketShaper before read/write access is granted.
Traffic Class	See **Class**.
Traffic Discovery	The PacketShaper process of observing and creating traffic classes for all packets as they pass through the unit. This process compiles a list of the protocols and applications in use on a network, creating a *traffic tree*.
Traffic Tree	A hierarchical list of traffic classes that are used on a network. The traffic discovery process creates the tree automatically.
Transaction	An exchange of packets between two hosts on a computer network.
Transaction Delay	The delay experienced by an application communicating between two hosts.
Transmission Control Protocol (TCP)	A reliable transport-layer method for managing end-to-end data flow between network-connected computers. The session-oriented component of TCP/IP.
Type-Of-Service (TOS)	A three-bit portion of the IP header in a TCP/IP packet reserved for controlling Quality of Service (QoS).
UDP	See User Datagram Protocol.

URL Category	A type of web traffic, such as Social Networking, Gambling, Pornography, News Media, and Shopping. With URL category classification, the PacketShaper examines the URL, determines what category it belongs to, and then classifies the traffic into the appropriate category class.
User Datagram Protocol (UDP)	An OSI model layer 4 network protocol that supports connectionless services, such as the IP suite.
UTC	An international time standard similar to Greenwich mean time. UTC stands for "coordinated universal time." When scheduling policies to run at a specific time, you can enter the time in UTC or local time.
Violating Host List	The host list to which adaptive response automatically adds the IP addresses of hosts that violate the agent's thresholds. See Agent Host Lists.
VLAN (Virtual LAN)	A local area network that maps workstations on some other basis than geographic location (for example, by department, type of user, or application).
WAFS	See Wide Area File Services.
WAN	See Wide Area Network.
Watch Mode	A non-inline, monitor-only mode.
WebPulse	A cloud-based service provided by Symantec that receives URL requests and responds with the category, web application, and/or operation associated with the URL. See WebPulse Overview.
Weighted Fair Queuing	A method of sequencing packets that defines bandwidth allocation according to an application's needs.
Wide Area File Services (WAFS)	A storage technology that allows organizations to access remote data centers as if they were local.
Wide Area Network (WAN)	One or more electronic circuits that connect computers in different locations.

Tutorials

PacketGuide tutorials go beyond merely describing different ways to configure your PacketShaper; they actually show you how to do it. Animated sequences walk you through each step, showing you exactly how to navigate through the browser interface to perform each task. Each tutorial also contains hints and suggestions from the pages of PacketGuide to help you along the way. Click on any of the tutorial titles to open that tutorial in a new browser window.

> **Note:** To view these tutorials, you must have a web browser with a Adobe Flash plug-in. If you do not have the Flash plug-in, you can download it from Adobe by clicking here.

- Tutorial: Web Application Classification and Reporting
  Learn how web application classification and reporting works on PacketShaper. For additional information, refer to the following topic: Web Application Classification Overview.

- Tutorial: User Classification and Reporting
  Classify traffic by Active Directory user names and view reports of top users. For additional information, refer to the following topic: User Awareness Overview.

- Tutorial: Detecting and Controlling Virus-Infested Hosts
  Detect problems on your network with the Host Analysis tool, and then take steps to correct the problems caused by virus-infested hosts. For additional information, refer to the following topic: *Detect and Limit DoS Attacks*.

- Tutorial: Managing Citrix Performance Detailed information to maximize the performance of all your Citrix applications. For additional information, refer to the following topic: *Manage Citrix Performance*.

- Tutorial: Managing Voice over IP (VoIP) Traffic
  Create VoIP classes and folders to protect VoIP performance, and ensure other applications do not get squeezed for bandwidth when VoIP demand increases. For additional information, refer to the following topic: *Manage Voice and Video Sessions*.

PacketShaper Feature Matrix

The following table lists major PacketShaper features and indicates the version each feature was introduced or removed.

Feature	Version Introduced
Captive Portal	11.8.1
enableJumboFrame system variable introduced	11.7.1
padVlanPkt system variable introduced	11.6.2
High availability: standby and access link monitoring	11.6.1
Customer portal administration	11.6.1
Email custom reports	11.6.1
Import new SSL certificate ("setup ssl install" on page 1345)	11.5.3
Regenerate self-signed SSL certificate (see "Recovery Feature" on page 327)	11.5.3
PolicyCenter support	11.5.2
Securing the inside/outside interfaces	11.5.2
RADIUS	11.5.1
TACACS+	11.5.1
Adaptive Response	11.5.1
Events	11.5.1
Email	11.5.1
Response Time Measurement	11.4.1
WebPulse Express on the PS-S500	11.3.1
Restriction that the PacketShaper must be managed from the MGMT port has been removed. (It is still recommended that you manage the PacketShaper via a separate management network, though.) This release includes a new and improved method for management of PacketShaper via the inline NICs as well as via the SYS MGMT port.	11.3.1
arp CLI command, with support for S-series NICs	11.3.1
cat CLI command, with new option for specifying number of lines to display	11.2.1.8
Support for PS-S200 and PS-S400	11.2.1
Password recovery	11.2.1
SNMP support	11.2.1
IntelligenceCenter support for PS-S500	11.1.14
Support for PS-S500	11.1.1
New licensing and image management infrastructure	11.1.1
Permanent removal of the following features: Xpress compression and acceleration Offline reports WCCP redirection Frame Relay and ATM Plug-ins Traffic information reporting CLI commands for file management	11.1.1

Feature	Version Introduced
Temporary removal of the following features: PolicyCenter and IntelligenceCenter support High Availability, including Direct Standby and Access Link Monitoring Failover setup and policy Site router configuration MPLS classification and tagging Response Time Measurement (all related ME variables and graphs) RADIUS and TACACS authentication Adaptive Response Events Email notification Support for SNMP MIBs, OIDs, and trap messages Using 207.78.98.254 for initial configuration (Guided Setup) Backup/restore of measurement data HTTPS, SSH settings; SSL cipher strength Securing inside/outside interfaces Data form generation	11.1.1
User group support	9.2.2
User awareness	9.2.1
Web application classification	9.2.1
Traffic information reporting	9.2.1
IPv6 sub-classification	9.1.1
Support for PolicyCenter on Windows Server 2008	8.7.2
Model class tree	8.7.1
Disable web applications	8.7.1
Configurable standby interface on PS12000	8.7.1
Support for 10,000 classes on PS12000 ISP	8.7.1
Display traffic statistics	8.7.1
PS 12000 support	8.6.3
SSL cipher strength	8.6.3
URL categories	8.6.1
Sky Host List Editor	8.6.1
Sky interface	8.5.1
Service groups	8.5.1
Heartbeats	8.5.1
User-defined services	8.4.1
SNMP3	8.3.1
TACACS+	8.3.1
WCCP	8.2.0
VoIP enhancements	8.2.0
Capture packets by NIC	8.2.0
Host Analysis tool	8.2.0
Statistical Graphing tool	8.2.0
Xpress	8.0.0

Product Information

Product-specific information about PacketShaper S-series

First Steps	Become productive with your PacketShaper almost immediately by following the instructions in this PDF file. These step-by-step instructions provide novices with the first steps to learning PacketShaper software.
Specifications	Dimensions, weight, number and type of ports, power requirements, and environmental specifications
Configuration Limits	Maximum values for PacketShaper 11.9 features
Power	Power-on and power-off instructions

Log in to PacketShaper Browser Interface

To log in to PacketShaper using a web browser:

1. Open a supported web browser. See "Browser Settings" on page 1028.

2. In the browser's address bar, type **https://*<IP_address>*** where *<IP_address>* is the IP address assigned to the PacketShaper. The login screen appears.

> **Caution:** You must log in to the PacketShaper using HTTPS, not HTTP. You will get an ERR_CONNECTION_TIMED_OUT message, such as *web page not available* or *this page cannot be displayed* message if you use HTTP.

3. Enter the User Name (**touch** or **look**).

4. Enter the Password.

5. Click **Login**. The the Sky UI user interface appears. See " Sky UI Overview " on page 106.

6. If you want to switch to the legacy user interface, click **Advanced UI**.

Log Out of PacketShaper

For security purposes, users should explicitly log out of PacketShaper:

- To log out of the browser interface, click the **LOG OUT** link in the banner or the **logout** button on the Advanced UI Info tab.
- To log out of the command-line interface, type **exit**.

Logging out discards session content and displays the login page. If a user doesn't explicitly log out, PacketShaper will automatically time out after one hour of inactivity. When a remote login session times out, PacketShaper sends a "timed out" message and disconnects.

> **Note:** When the Sky user interface has been loaded, the real-time graphs are constantly polling the PacketShaper so the session does not time out.

You can also use the **logout** button/link to change the access level. For example, if you are currently logged in as a touch user, you can log out and then log back in with the look user name and password.

First Steps to Using...

The following PDF contains tutorials for getting you started using PacketShaper. Follow the step-by-step instructions to learn how to classify network traffic, analyze network and application performance, and solve performance problems.

- First Steps to Using a PacketShaper

Additional PDFs that document how to use specific PacketShaper classification features are also available:

- User Classification
- URL Categories
- Web Application Classification
- IPv6 Classification

PacketShaper S-Series Product Specifications

Look up the dimensions, weight, number and type of ports, power requirements, and environmental specifications.

- PS-S200
- PS-S400
- PS-S500

PS-S400 Product Specifications

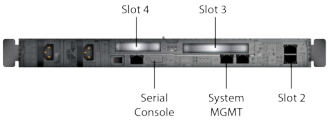

See also "PS-S400 Device and Port Names" on page 942.

Dimensions	Standard 1U 19-inch rack mount Height: 1.69 in. (4.29 cm) Width: 17.00 in. (43.25 cm) Depth: 22.52 in. (57.2 cm)
Weight	33 lbs. (15 kg)
Built-In NICs	1 Gb bypass
Expansion Slots	2
System Management Port	Yes (SYS MGMT) **Note:** The BMC MGMT port is not for customer use.
Serial Port	RS-232 (AT-compatible) serial port with male DB-9 connector
Power LED	Not lit: Appliance is powered off . Amber: Appliance has been powered on and is in the process of booting. Green: Appliance is powered on and has finished the boot process.
RJ-45 Status LEDs	**Link Status (Upper/Left) LED:** Illuminated green when the network cable is properly connected on both ends. Flickers when the unit is transmitting and receiving data. **Speed Status (Lower/Right) LED:** Yellow: 1 Gbps link speed Green: 10 Gbps link speed Blinking green: 10 Gbps link speed with bypass Blinking yellow: Disconnect
Electrical Data	Maximum Power: AC 450 watts, DC 770 watts Power Input: AC or DC power Input Voltage Rating: AC: 100-127/200-240V~, DC: -48VDC (+/= 10% tolerance) Input Current Rating: AC 8A/4A, DC 24A Max Input Frequency Rating: 47 Hz ~63 Hz Output Power: Voltage: 12VDC Current: AC 37A, DC 62.5A Ripple: 120mV Total Regulation: +-0.5% Maximum PCIe Power Load: 25 watts
Power Supplies	1 or 2 (a second power supply can be purchased as an option)
Fans	6

Drives	2 x HDD: 500GB - 2.5" SATA 1 x SSD: 8GB Slim SATA Module
Memory	2 x DDR3 (8GB), 16GB total
Environmental (Operating)	Temperature: 41° F to 104° F (5° C to 40° C) at sea level 41° F to 95° F (5° C to 35° C) at 10,000 ft. Relative humidity: 20% to 85% non-condensing Altitude: 0 to 10,000 ft. (3000 M) operating
Cooling Flow	45 CFM Front to back cooling method
Heat Dissipation	1512 BTU/hour
Vibration	0.17 G's, 5 - 500Hz, swept-sine Sweep Speed: 1 octave/min Duration: Three sweeps on each axis (X, Y, Z)
Half-Sine Shock	3.5 G's, 11 ms, half-sine
Managed Link Speed	Multiple configurations are offered; supports links up to 10 Gbps full duplex. Software configuration determines maximum shaping capacity.
RoHS Compliant	Yes

Note: Specifications are subject to change without notice.

Warning Class 1 laser product
Attention Produit laser de classe 1
Warnung Laserprodukt der Klasse 1

PS-S400 Device and Port Names

The PS-S400 can have dual- or quad-port network interface cards (NICs) installed in Slots 3 and 4. The slots and ports are numbered as illustrated in the diagram below.

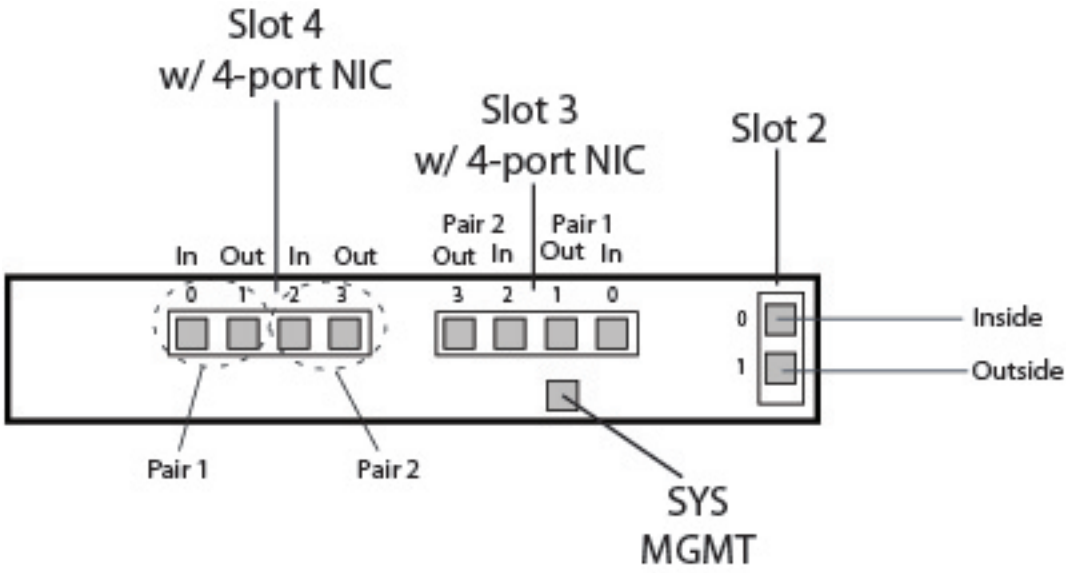

Device Name	Card Type	Port Type	Pair Position	Port Number	Port Name
Slot2	built-in	Inside		0	slot2_in
		Outside		1	slot2_out
Slot3	2-port	Inside		0	slot3_in
		Outside		1	slot3_out
Slot3_pair1	4-port	Inside	Right	0	slot3_in1
		Outside	Right	1	slot3_out1
Slot3_pair2		Inside	Left	2	slot3_in2
		Outside	Left	3	slot3_out2
Slot4	2-port	Inside		0	slot4_in
		Outside		1	slot4_out
Slot4_pair1	4-port	Inside	Left	0	slot4_in1
		Outside	Left	1	slot4_out1
Slot4_pair2		Inside	Right	2	slot4_in2
		Outside	Right	3	slot4_out2

When configuring the PS-S400, some operations require the specification of device names or port names. The table above provides the proper names to use.

PS-S400 NIC Support

The PS-S400 can have dual- or quad-port network interface cards (NICs) installed in Slots 3 and 4. The slots are numbered as illustrated in the diagram below.

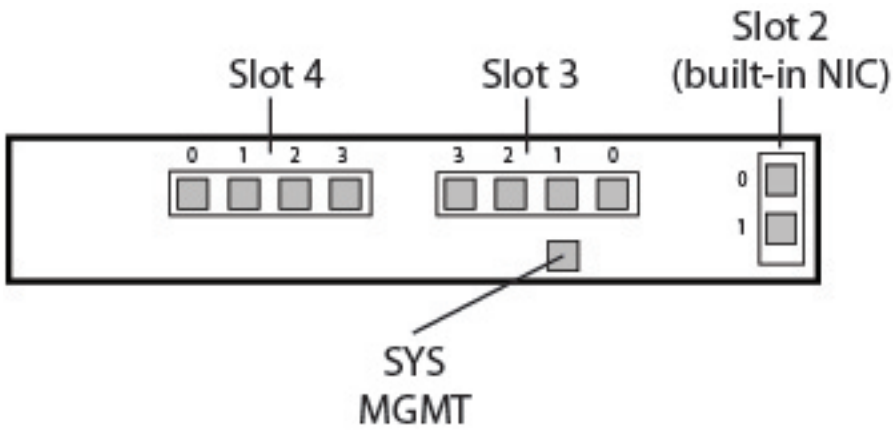

The table below indicates which NICs are supported in each slot. (✓=supported)

Slot #	Length	Quad 1G Bypass Copper NIC-S400-4xGigE-PT	Quad 1G Bypass Fiber SX NIC-S400-4xGigE-PT-Fiber	Quad 1G Bypass Fiber LX NIC-S400-4xGigE-LX-Fiber	Dual 10G Bypass Copper NIC-S400-2x10G-PT	Dual 10G Bypass Fiber SR NIC-S400-2x10G-PT-SR-Fbr	Dual 10G Bypass Fiber LR NIC-S400-2x10G-PT-LR-Fbr
Slot 3	Full-height	✓	✓	✓	✓	✓	✓
Slot 4	Half-height	✓	NA	NA	✓	✓	✓

Note: All fiber-optic NICs have LC small form-factor connectors.

Caution: When you add a new NIC, remove an existing NIC, or move a NIC to a different slot, the NIC speed and duplex settings will be reset to their defaults. After making a NIC change, make sure to verify the NIC settings and modify to desired values if necessary.

PS-S500 Product Specifications

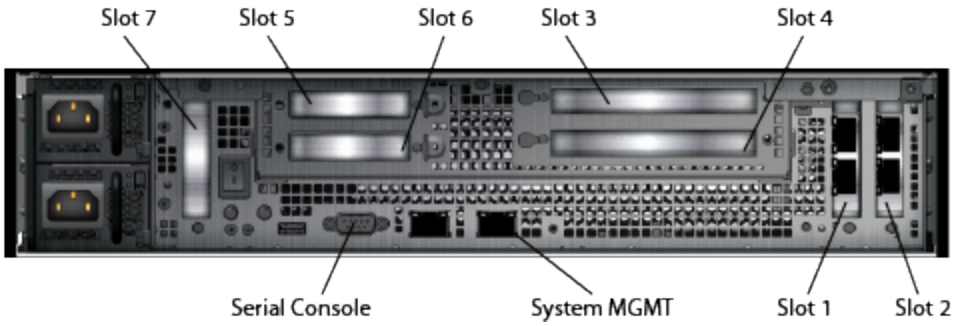

See also "PS-S500 Device and Port Names" on page 946.

Dimensions	Standard 2U 19-inch rack mount Height: 3.43 in. (8.72 cm) Width: 17.06 in. (43.33 cm) Depth: 27.95 in. (71.0 cm) excluding bezel
Weight	66 lbs. (30 kg)
Built-In NICs	No bypass-capable pairs are on the base system. (The two 1 GbE NICs on the base system are for system management of the unit.) 4 (Slots 1, 2, 3, 4) See "PS-S500 NIC Support " on page 947.
Expansion Slots	**Note:** Slots 5-7 cannot currently be used on the PS-S500.
System Management Port	Yes (SYS MGMT) **Note:** The BMC MGMT port is not for customer use.
Serial Port	RS-232 (AT-compatible) serial port with male DB-9 connector
Power LED	Not lit: appliance is powered off Green: appliance is powered on
RJ-45 Status LEDs	**Link Status (Upper/Left) LED:** Illuminated green when the network cable is properly connected on both ends. Flickers when the unit is transmitting and receiving data. **Speed Status (Lower/Right) LED:** Yellow: 1 Gbps link speed Green: 10 Gbps link speed Blinking green: 10 Gbps link speed with bypass Blinking yellow: Disconnect

Electrical Data	Maximum Power: 1100 watts Power Input: AC or DC power Input Voltage Rating: 100-240V~ or 48-60VDC Input Current Rating: 12-5A, or 30A Input Frequency Rating: 50-60 Hz Output Power: DC power Voltage: 12VDC Current: Min 78A, Max 95A Ripple: 250mV Total Regulation: +-0.5% Maximum PCIe Power Load: 25 watts
Power Supplies	2
Fans	5
Drives	2 x HDD: 1TB - 2.5" NL SAS 2 x SSD: 8GB Slim SATA Module
Memory	8 x 8GB DDR3 DIMMs, 64GB total
Heat Dissipation	Typical: 2598.42 BTU/hour, Max: 3751 BTU/hour
Environmental (Operating)	Temperature: 41° F to 104° F (5° C to 40° C) at sea level 41° F to 95° F (5° C to 35° C) at 10,000 ft. Relative humidity: 20% to 85% non-condensing Altitude: 0 to 10,000 ft. (3000 M) operating
Cooling Flow	107 CFM Front to back cooling method
Heat Dissipation	3593 BTU/hour
Vibration	0.17 G's, 5 - 500Hz, swept-sine Sweep Speed: 1 octave/min Duration: Three sweeps on each axis (X, Y, Z)
Half-Sine Shock	3.5 G's, 11 ms, half-sine
Managed Link Speed	Multiple configurations are offered; supports links up to 10 Gbps full duplex. Software configuration determines maximum shaping capacity.
RoHS Compliant	Yes

Note: Specifications are subject to change without notice.

Warning Class 1 laser product
Attention Produit laser de classe 1
Warnung Laserprodukt der Klasse 1

PS-S500 Device and Port Names

The PS-S500 can have up to four network interface cards (NICs) installed in slots 1 through 4. (Slots 5-7 are not currently available.) Slots 3 and 4 support dual-port and quad-port NICs. The slots and ports are numbered as illustrated in the diagram below.

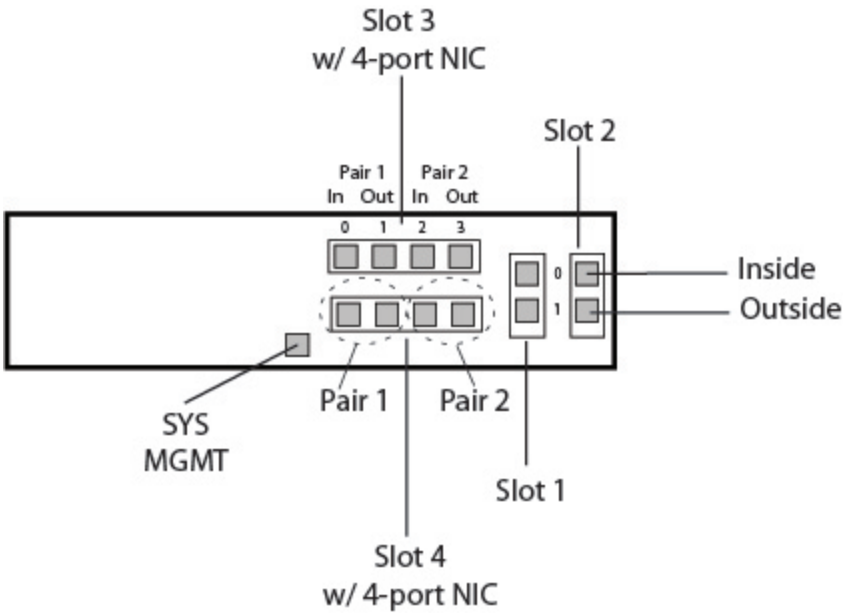

Device Name	Card Type	Port Type	Pair Position	Port Number	Port Name
Slot1	2-port	Inside		0	slot1_in
		Outside		1	slot1_out
Slot2	2-port	Inside		0	slot2_in
		Outside		1	slot2_out
Slot3	2-port	Inside		0	slot3_in
		Outside		1	slot3_out
Slot3_pair1	4-port	Inside	Left	0	slot3_in1
		Outside	Left	1	slot3_out1
Slot3_pair2		Inside	Right	2	slot3_in2
		Outside	Right	3	slot3_out2
Slot4	2-port	Inside		0	slot4_in
		Outside		1	slot4_out
Slot4_pair1	4-port	Inside	Left	0	slot4_in1
		Outside	Left	1	slot4_out1
Slot4_pair2		Inside	Right	2	slot4_in2
		Outside	Right	3	slot4_out2

When configuring the PS-S500, some operations require the specification of device names or port names. The table above provides the proper names to use.

PS-S500 NIC Support

The PS-S500 can have up to four network interface cards (NICs) installed in slots 1 through 4. (Slots 5-7 are not currently available.) Slots 3 and 4 support dual-port and quad-port NICs. The slots are numbered as illustrated in the diagram below.

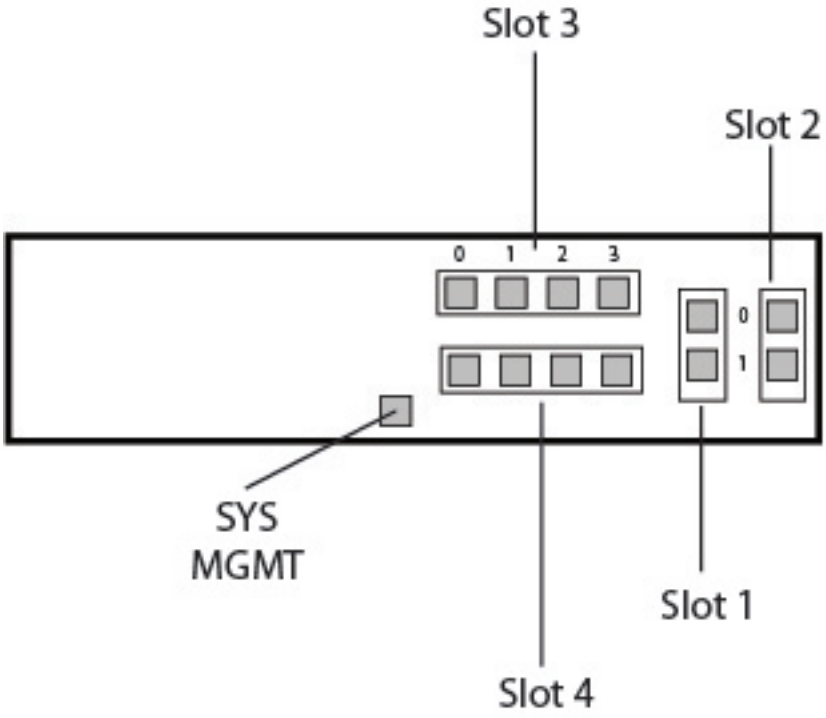

The table below indicates which NICs are supported in each slot. (✓=supported)

Slot #	Length	Quad 1G Bypass Copper NIC-S500-4xGigE-PT	Quad 1G Bypass Fiber SX NIC-S500-4xGigE-PT-Fiber	Quad 1G Bypass Fiber LX NIC-S500-4xGigE-LX-Fiber	Dual 10G Bypass Copper NIC-S500-2x10G-PT	Dual 10G Bypass Fiber SR * NIC-S500-2x10G-PT-SR-Fbr	Dual 10G Bypass Fiber LR NIC-S500-2x10G-PT-LR-Fbr
Slot 1	Half-height	✓	NA	NA	✓	✓	✓
Slot 2	Half-height	✓	NA	NA	✓	✓	✓
Slot 3	Full-height	✓	✓	✓	✓	✓	✓
Slot 4	Full-height	✓	✓	✓	✓	✓	✓

* PS-S500, PS-S400, and PS12000 appliances all use the same 10G fiber-optic SR NIC. Out of the box, the NIC has an S400 low-profile bracket attached to it. You will need to remove this bracket and attach the S500 low-profile bracket included in the box (PN 425-02875). The other bracket in the box (PN 425-02815) does not have a slot in it and works with the PS12000.

> **Note:** All fiber-optic NICs have LC small form-factor connectors.

Caution: When you add a new NIC, remove an existing NIC, or move a NIC to a different slot, the NIC speed and duplex settings will be reset to their defaults. After making a NIC change, make sure to verify the NIC settings and modify to desired values if necessary.

PacketShaper S-Series Network Interface Capabilities

The supported link rates and duplex modes for each available network interface on the PacketShaper S-Series are listed below.

MGMT port

10-half duplex
10-full duplex
100-half duplex
100-full duplex
1000-full duplex

All speeds below 1000 support both manual and auto-negotiation.
At 1000 bps, the device supports only auto-negotiation.

4-port 1G Ethernet NIC

10-half duplex
10-full duplex
100-half duplex
100-full duplex
1000-full duplex

The device supports manual and auto-negotiation for all supported speeds.

4-port 1G fiber-optic NIC

1000-full duplex

2-port 10G fiber-optic NIC(SR and LR)

10000-full duplex

This device supports only a manual 10000-full duplex configuration.

2-port 10G Ethernet NIC

100-full duplex
1000-full duplex
10000-full duplex

This device supports only auto-negotiation for all supported speeds.

> **Note:** 10G NICs are not supported on PS-S200.

PacketShaper S-Series Configuration Limits

Look up the maximum values for PacketShaper classes, partitions, policies, and other components.

"PacketShaper S200 Configuration Limits" on the facing page

"PacketShaper S400 Configuration Limits" on page 953

"PacketShaper S500 Configuration Limits" on page 954

Related Information:

- Currently shipping PacketShaper models and configurations
- End of Sale / End of Life Documents

PacketShaper S200 Configuration Limits

The following table lists maximum values for common components in PacketShaper 11.9 for the PacketShaper S200 models.

Component	PS-200-10ML	PS-200-50ML	PS-200-100ML*	PS-200-100MH*	PS-S200-250MH	PS-S200-500MH
Supported WAN link rates (full duplex)	10 Mbps	50 Mbps	100 Mbps	100 Mbps	250 Mbps	500 Mbps
Concurrent flows (TCP / other IP) **	5,000/ 2,500	100,000/ 50,000	100,000/ 50,000	200,000/ 100,00	275,000/ 150,000	350,000/ 200,000
Optimal Max # of IP hosts at one time *	7,500	25,000	25,000	75,000	200,000	350,000
Classes	512	1,024	1,024	2,000	3,000	5,000
Matching rules	2,000	5,000	5,000	5,000	7,500	10,000
Matching rules per class	250	1,000	1,000	1,000	1,000	1,000
Policies	512	1,024	1,024	2,000	3,000	5,000
Static partitions	512	1,024	1,024	2,000	3,000	5,000
Dynamic partitions	1,000	2,500	2,500	5,000	7,500	10,000
Network interface connection speeds	10 Mbps, 100 Mbps, 1 Gbps auto-detect					
Copper NIC options	4-port 10/100/1000BASE-T					
Fiber-Optic NIC options	4-port 1GbE SR					
Maximum URL category cache size (entries and bytes)	500,000 22 MB	500,000 22 MB	500,000 22 MB	1,000,000 44 MB	1,500,000 66 MB	2,000,000 88 MB
Maximum number of Active Directory users/groups+	100/50	200/100	200/100	2500/200	5000/225	10,000/250

Additional Information

Not all capacity specifications can be maximized simultaneously.

L = Low User Counts and H = High User Counts

*The PS-S200 offers models with high (H) and low (L) user counts, for example, PS-200-100ML and PS-200-100MH. The "L" models support fewer classes, partitions, dynamic partitions, and concurrent connections than the "H" models, since systems with fewer users have lighter requirements.

** These numbers have been rounded. PacketShaper can support more hosts and flows, however these figures represent the ideal maximums for producing optimal results. It is normal and expected to see the number of concurrent hosts at its limit.

+ The maximum number of groups is the number of unique groups that can be managed and reported on a PacketShaper. Although your Active Directory can have up to 1024 groups, PacketShaper can only manage and report on a fixed number of groups. When you create a group-based class or add a group to a user list, that group becomes a *group of interest* and is added to an internal groups of interest list. When you delete a group-based class or remove a group from a user list, that group is removed from the groups of interest list.

PacketShaper S400 Configuration Limits

The following table lists maximum values for common components in PacketShaper 11.9 for the PacketShaper S400 models.

Component	PS-S400-250MH	PS-S400-500MH	PS-S400-1GH	PS-S400-2GH
Supported WAN link rates (full duplex)	250 Mbps	500 Mbps	1 Gbps	2 Gbps
Concurrent flows (TCP / other IP) *	275,000/ 150,000	350,000/ 200,00	1,000,000/ 500,000	1,000,0000/ 500,00
Optimal Max # of IP hosts at one time *	200,000	350,000	450,000	550,000
Classes	3,000	5,000	7,500	10,000
Matching rules	7,500	10,000	15,000	25,000
Matching rules per class	1,000	1,000	1,000	1,000
Policies	3,000	5,000	7,500	10,000
Static partitions	3,000	5,000	5,000	10,000
Dynamic partitions	7,500	10,000	15,000	20,000
Network interface connection speeds		10 Mbps, 100 Mbps, 1 Gbps auto-detect, 10 Gbps		
Copper NIC options		4-port 10/100/1000BASE-T 2-port 10GBASE-T		
Fiber-Optic NIC options		4-port 1GbE SX 2-port 10Gbps SR 2-port 10Gbps LR		
Maximum URL category cache size (entries and bytes)	1,500,000 66 MB	2,000,000 88 MB	2,000,000 88 MB	2,000,000 88 MB
Maximum number of Active Directory users/groups+	5000/225	10,000/250	10,000/250	10,000/250

Additional Information

Not all capacity specifications can be maximized simultaneously.

* These numbers have been rounded. PacketShaper can support more hosts and flows, however these figures represent the ideal maximums for producing optimal results. It is normal and expected to see the number of concurrent hosts at its limit.

+ The maximum number of groups is the number of unique groups that can be managed and reported on a PacketShaper. Although your Active Directory can have up to 1024 groups, PacketShaper can only manage and report on a fixed number of groups. When you create a group-based class or add a group to a user list, that group becomes a *group of interest* and is added to an internal groups of interest list. When you delete a group-based class or remove a group from a user list, that group is removed from the groups of interest list.

PacketShaper S500 Configuration Limits

The following table lists maximum values for common components in PacketShaper 11.9 for the PacketShaper S500 models.

Component	PS-S500-2GH	PS-S500-5GH	PS-S500-10GH
Supported WAN link rates (full duplex)	2 Gbps	5 Gbps	10 Gbps
Concurrent flows (TCP / other IP) *	1,000,000/ 500,000	1,500,000/ 1,000,000	2,000,000/ 2,000,000
Optimal Max # of IP hosts at one time *	550,000	575,000	600,000
Dynamic partitions	20,000	35,000	55,000
Static partitions		10,000	
Classes		10,000	
Policies		10,000	
Matching rules		25,000	
Matching rules per class		1,000	
Network interface connection speeds		10 Mbps, 100 Mbps, 1 Gbps auto-detect, 10 Gbps	
Copper NIC options		4-port 10/100/1000BASE-T 2-port 10GBASE-T	
Fiber-Optic NIC options		4-port 1GbE SR 2-port 10Gbps SR 2-port 10Gbps LR	
Maximum URL category cache size (entries and bytes)		2,000,000 88 MB	
Maximum number of Active Directory users/groups+		10,000/250	

Additional Information

Not all capacity specifications can be maximized simultaneously.

* These numbers have been rounded. PacketShaper can support more hosts and flows, however these figures represent the ideal maximums for producing optimal results. It is normal and expected to see the number of concurrent hosts at its limit.

+ The maximum number of groups is the number of unique groups that can be managed and reported on a PacketShaper. Although your Active Directory can have up to 1024 groups, PacketShaper can only manage and report on a fixed number of groups. When you create a group-based class or add a group to a user list, that group becomes a *group of interest* and is added to an internal groups of interest list. When you delete a group-based class or remove a group from a user list, that group is removed from the groups of interest list.

Power

For instructions on powering the PacketShaper S-Series on and off, select the model from the list below:

- PS-S200
- PS-S400
- PS-S500

PS-S200 Power

The PS-S200 has a single power supply. Symantec recommends connecting the PacketShaper to an uninterruptable power supply (UPS) to avoid configuration corruption that can occur if the appliance loses power unexpectedly.

The power switch on the back panel of the appliance should be used for power cycling the unit, not for powering down. This switch does not completely power down the appliance; when the PS-S200 is turned off with this switch, you can no longer access any of the user interfaces or ping the unit, but there is still power going to the appliance.

Powering On

1. Connect the power cord to the PacketShaper's power socket in the back of the unit.
2. Plug the other end of the power cord into AC power.
3. If the appliance does not automatically power on, switch the appliance on by pressing the rear soft power switch.

Power Cycling

To power cycle (reset) the PS-S200:

1. Press and release the rear soft power switch.
2. Wait a moment.
3. Press and release the power switch again.

Powering Off

1. Go to the command-line interface.
2. Type shutdown and press Enter.
3. Unplug the power cord(s).

Caution: Do NOT turn off the appliance power without executing the proper shutdown procedure. Unplugging the power cord when the appliance is turned on will corrupt the file system.

PS-S400 Power

The PS-S400 ships with a single power supply and a second one can be purchased, if desired. When two power supplies are installed, they are alternating hot-swappable. For information about replacing power supply modules, see *S-Series Maintenance and Upgrade Guide*.

Two power cables are included with your PacketShaper – one for each power supply. Ideally, each power cable should be connected to a different power source. In addition, Symantec recommends connecting the PacketShaper to an uninterruptable power supply (UPS) to avoid configuration corruption that can occur if the appliance loses power unexpectedly.

The power switch on the back panel of the appliance should be used for power cycling the unit, not for powering down. This switch does not completely power down the appliance; when the PS-S400 is turned off with this switch, you can no longer access any of the user interfaces or ping the unit, but there is still power going to the appliance.

Powering On

1. Connect the power cord(s) to the PacketShaper's power sockets in the back of the unit.
2. Plug the other end of the power cord(s) into AC power. When two power supplies are installed, be sure to connect the two power cords to outlets on separate circuit breakers.
3. If the appliance does not automatically power on, switch the appliance on by pressing the rear soft power switch.

Power Cycling

To power cycle (reset) the PS-S400:

1. Press and release the rear soft power switch.
2. Wait a moment.
3. Press and release the power switch again.

Powering Off

1. Go to the command-line interface.
2. Type shutdown and press Enter.
3. Unplug the power cord(s).

Caution: Do NOT turn off the appliance power without executing the proper shutdown procedure. Unplugging the power cords when the appliance is turned on will corrupt the file system.

PS-S500 Power

The PS-S500 has dual hot-swappable alternating power supplies. For information about replacing power supply modules, see *S-Series Maintenance and Upgrade Guide*.

Two power cables are included with your PacketShaper – one for each power supply. Ideally, each power cable should be connected to a different power source. In addition, Symantec recommends connecting the PacketShaper to an uninterruptable power supply (UPS) to avoid configuration corruption that can occur if the appliance loses power unexpectedly.

The power switch on the back panel of the appliance should be used for power cycling the unit, not for powering down. This switch does not completely power down the appliance; when the PS-S500 is turned off with this switch, you can no longer access any of the user interfaces or ping the unit, but there is still power going to the appliance.

Powering On

1. Connect the power cords to the PacketShaper's power sockets in the back of the unit.
2. Plug the other ends of the power cords into AC power. Be sure to connect the two power cords to outlets on separate circuit breakers.
3. If the appliance does not automatically power on, switch the appliance on by pressing the rear soft power switch.

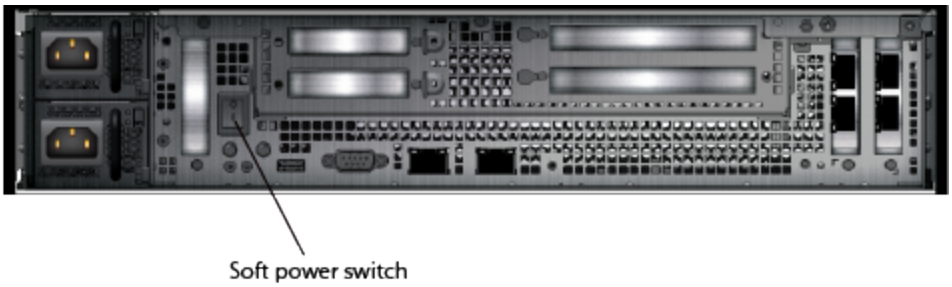

Power Cycling

To power cycle (reset) the PS-S500:

1. Press and release the rear soft power switch.
2. Wait a moment.
3. Press and release the power switch again.

Powering Off

1. Go to the command-line interface.
2. Type shutdown and press Enter.
3. Unplug the power cord(s).

Caution: Do NOT turn off the appliance power without executing the proper shutdown procedure. Unplugging the power cords when the appliance is turned on will corrupt the file system.

Deployment Topologies

Getting the most from a PacketShaper depends on choosing the right deployment strategy for your topology and needs. That's what these deployment topologies are all about: where, when, and what to deploy.

The following deployment topologies are relevant after you decide you want PacketShapers but before you start unplugging and plugging cables to install them. Although these sections of PacketGuide address you, the customer, it assumes that a Symantec sales engineer or partner is working with you in deciding on an appropriate deployment strategy.

You can deploy PacketShaper comprehensively throughout many or all offices, or you can start more slowly and follow a phased deployment strategy starting with a core site first. Benefits vary with the product and locations you choose.

> **Note:** For further assistance in deploying the PacketShaper, see Deployment Assistance.

Select a link below to view deployment and configuration information for the following topics:

> **Note:** For best performance, Symantec recommends managing the PacketShaper appliance via a separate management network. However, you can also manage the PacketShaper via inline NICs if necessary. To enforce all management traffic for PacketShaper S-series to its management port only, see KB article 000032462.

Description	Topology
"Core Site WAN Link Topology" on page 972: A core-site WAN-link topology has a PacketShaper only at core sites.	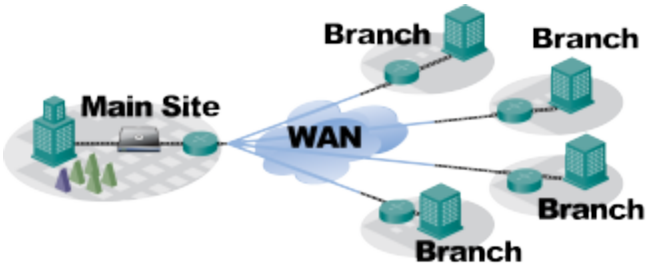
"Core Site Internet Link Topology" on page 969: A core-site Internet-link topology has a PacketShaper at the core site's link to VPN-connected branch offices, dial-up VPN users, partner extranets, and/or simply the web.	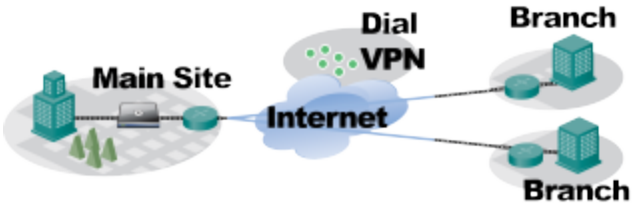

Description	Topology
"Distributed Branch Topology" on the facing page: A distributed topology has a PacketShaper at each core and remote site. It is the most powerful and flexible topology.	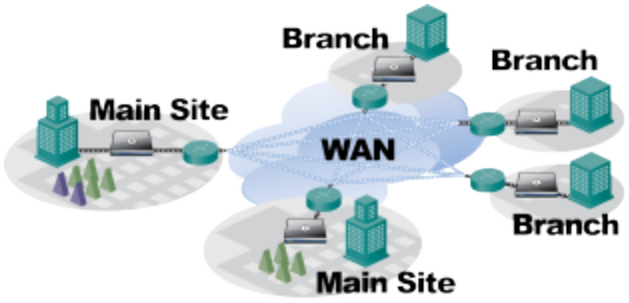
"Topologies with Multiple LANs " on page 974: This topology includes any environment with multiple LANs connected to one or more routers attached to WAN or Internet links.	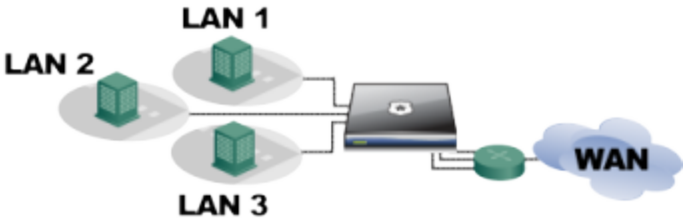
"Virtual Private Network Topology" on page 993: A VPN can serve as an economical WAN alternative, but may suffer unpredictable performance for mission-critical applications. To add to this problem, most monitoring tools aren't able to portray VPNs' encrypted traffic with any detail. A PacketShaper can help both issues.	
"High Availability Topologies" on page 977: PacketShaper's standby function allows a PacketShaper to integrate smoothly into a variety of redundant network topologies – those with redundant routers, redundant switches, two data paths, two LANs, and/or redundant firewalls.	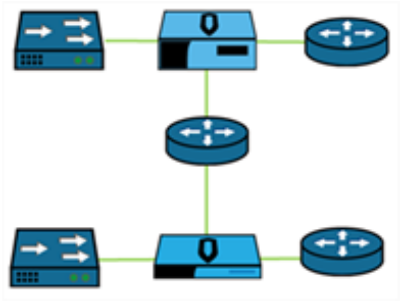
"Topologies with Firewalls" on page 967: This topology includes any environment with both a firewall and a PacketShaper.	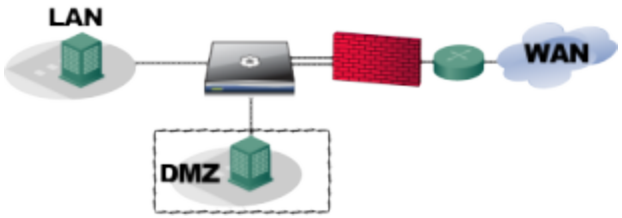
"Non-Inline Deployments (Watch Mode)" on page 996: PacketShaper's *watch mode* is used to monitor traffic non-inline, where the PacketShaper is not cabled into the core data path. This type of deployment is best suited for enterprises that have strict change control procedures or security restrictions on the introduction of inline elements into the network.	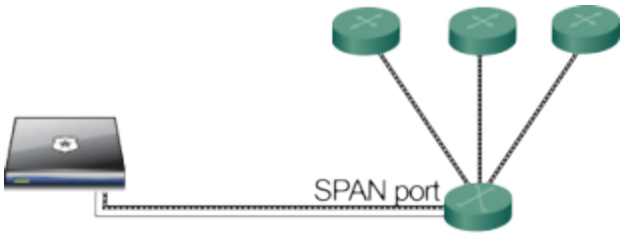

Distributed Branch Topology

A distributed topology has a PacketShaper at each core and remote site. It is the most powerful and flexible topology.

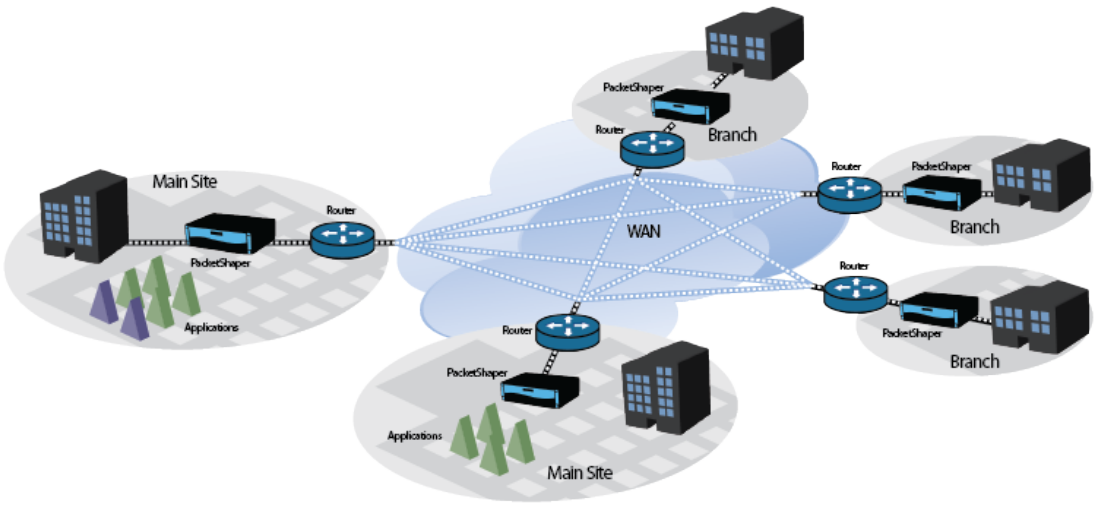

Environmental Requirements

Token ring networks do not support this topology, but there are no other special requirements; PacketShaper appliances are interoperable and fully compatible with each other.

Any core site or branch office can host applications, and there is no specific limit to the number of branch offices supported by this topology; ten to several hundred branch locations are typical. The greater the number of branch offices, the more Symantec's centralized-management products reduce the cost and overhead involved in a distributed solution.

Reasons to Choose

Customers choose this topology because it is the most powerful and flexible, and it scales well. Specific characteristics that indicate a distributed deployment is best include:

- Many branch offices need management
- A meshed network design (traffic doesn't pass through a core site)
- Traffic between branch offices (VoIP, shared folders, collaboration tools, and so on)
- Independent Internet links at branch offices (not dependent on a core site's Internet link)

- Applications hosted on servers at branch offices
- Web-based applications on high-latency links
- Small branch offices must access centralized applications through very constrained links
- Connectionless applications that must be contained. PacketShaper's TCP Rate Control can contain TCP traffic in both directions, but UDP and legacy flows require PacketShaper on both sides of a connection for optimal management of inbound flows.

Installation Considerations

Purchase of multiple appliances is required. Most installations require the same procedures as described in the unit's *Quick Start Guide*. Plug the router's LAN port to PacketShaper's OUTSIDE port and plug the LAN to PacketShaper's INSIDE port, then follow the Guided Setup prompts.

Capabilities

- A PacketShaper offers all its monitoring and analysis features, including classification, utilization analysis, service-level management, efficiency analysis, thresholding and events, packet capture, diagnostic graphs, and reports.
- When shaping is enabled, a PacketShaper also provides control features, including those to protect, contain, pace, and provision bandwidth, provide TCP Rate Control, mark packets, and suppress DoS attacks.

Limitations

A PacketShaper at only the core site does not aid performance at other branch offices with their own Internet links. It also does not aid performance for WAN traffic that uses a different link and network.

Configuration Suggestions

If you do not manage traffic to multiple branch offices on your Internet link, use a traffic tree that is based upon network applications, rather than branch locations or subnet. Automatic traffic discovery and classification are a good basis for the traffic classes, but you can also add or tailor classes based on other criteria. It's a simple tree to create and use, managing traffic according to type of traffic (such as application) rather than according to destination or location. Suggestions include:

- Customize an application-based tree with some of the web classification criteria to differentiate customers versus partners versus employees versus any other group you want to protect or control separately.

- Create a folder for streaming media and move all streaming protocols underneath it. Assign rate policies with guaranteed, per-session bandwidth. Consult PacketGuide's Recommendations for instructions.

- Create a folder called Unsanctioned or Undesirable or any other name that suits. Move any traffic classes that you would like to block (or at least contain to small amounts of bandwidth) to the folder. Then assign a policy to limit or block the traffic. Consult PacketGuide's Recommendations for instructions.

- If you use VPN, create a folder class to handle all VPN traffic types. Then assign a carefully sized partition to protect your VPN traffic. Use a dynamic subpartition on a VPN class if you want to provision bandwidth for each dial-up user.

- If you do manage multiple branch offices' traffic on your core site's Internet link, consider using instead the VPN topology.

Topology Example

Example

This example shows a very simple point-to-point topology.

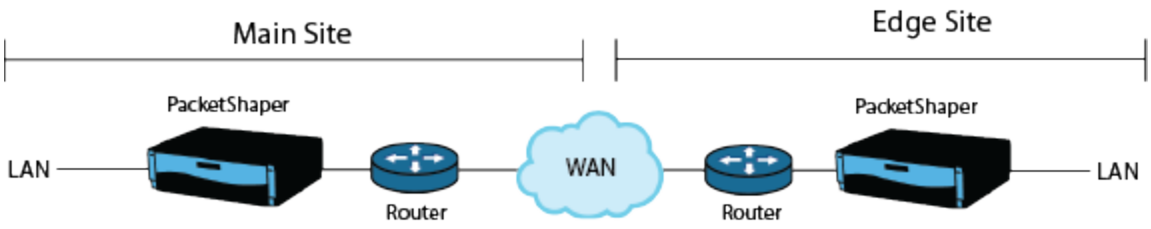

Scaling Considerations

This topology is most suited and adaptable to large or growing networks.
Each branch office needs the PacketShaper model that matches its WAN access-link capacity. Consult the Configuration Limits to check which model suits your tree's needs.

The type of traffic tree you choose for a core site impacts the model you select. Choices for your traffic tree and implications are discussed in *Traffic Tree Options*. The model for a core site needs to accommodate the link's capacity, the anticipated maximum number of concurrent flows, and the anticipated number of users. If you choose a location-based tree for the core site, the model needs to accommodate the number of traffic classes and matching rules required for each remote site in the tree.

Configuration Suggestions

For each branch office, use an application-based tree. Assign partitions and policies to each application that needs protection or that undermines other applications' performance. Consult Policy and Partition Guidelines for suggestions on specific applications and Control Strategy Overview for more high-level advice.

For the core site or sites, use either a simple location-based tree (to control only the total bandwidth going to each location), or a location-based tree with per-location applications (to control each application for each location separately).

- **If you are using a simple location-based tree,** assign a static partition to each location with a size equal to the branch's capacity, and leave it to the branch PacketShaper appliances to manage individual applications.

- **If using a location-based tree with per-location applications and** branch offices' application traffic goes through the core site, give each branch office a static partition with the proper minimum/maximum sizes corresponding to the branch WAN-link capacity (CIR or port speed). Then, for each branch office's applications, assign partitions and policies that correspond to those you assigned at the branch office's unit. UDP traffic benefits most from having policies at both ends. TCP Rate Control on your branch office PacketShapers manages both inbound and outbound rates for your TCP-based traffic classes.

Topologies with Firewalls

This topology includes any environment with both a firewall and a PacketShaper.

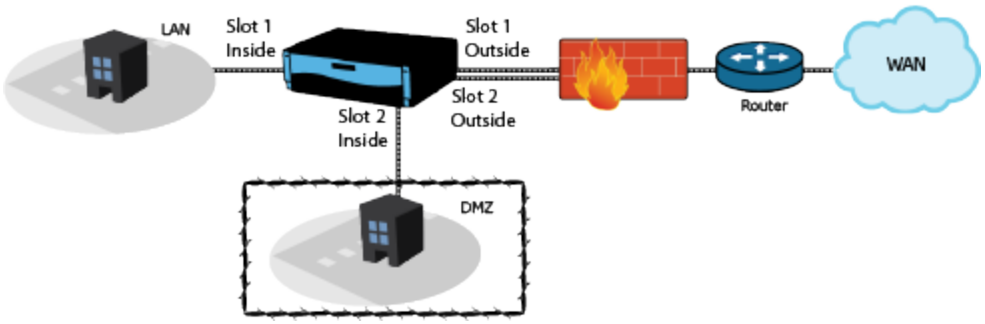

Installation Considerations

The issues for this topology center around installation and cabling of the router, firewall, and a PacketShaper. The additions of NAT (network address translation) and/or a DMZ (publicly available content, not protected by the firewall) require consideration.

Placing a PacketShaper appliance between the firewall and the protected LAN results in the most function and least confusion. If you have NAT, these appliances deal only with internal IP addresses. If your firewall is a termination point for VPN tunnels, then they see only unencrypted traffic. With the appliance between the firewall and unprotected LAN, you get full application-intelligent classification features and the ability to control individual types of traffic (instead of just one large mass of IPSec traffic). If you use the Top Listeners feature, then you get a meaningful list of the top consumers of any traffic class with their real, internal IP addresses.

It is possible to deploy a PacketShaper between the router and the firewall, but it is not recommended. The appliance would see only NAT configured addresses; all candidates for the Top Talkers list and Top Listeners would appear to be just one address, and VPN classification and control would be very coarse.

With no DMZ:

1. Plug the router's LAN port to the firewall.
2. Connect the firewall's LAN port to the appliance's OUTSIDE port.
3. Connect the protected LAN to the appliance's INSIDE port.

With a DMZ:

1. Install a NIC (network interface card) into the PacketShaper. It will handle the traffic to and from the DMZ.
2. Plug the router's LAN port to the firewall.

3. Connect the firewall's port for all DMZ traffic to the NIC's OUTSIDE port.

4. Connect the NIC's INSIDE port to the DMZ LAN.

5. Connect the firewall's port for the protected LAN traffic to appliance's OUTSIDE port.

6. Connect the protected LAN to appliance's INSIDE port.

Capabilities

A PacketShaper offers all its monitoring and analysis features with no NAT, and all but one (either Top Talkers or Top Listeners) with NAT.

When shaping is enabled, a PacketShaper also offers control features on all traffic that passes. For example, you can allocate one third of your link to the DMZ. If the DMZ is not in use, its bandwidth will be loaned to others.

Limitations

With NAT, neither Top Talkers nor Top Listeners is useful (see *Installation Considerations*, above).

Configuration Suggestions

Define a static partition for DMZ traffic to protect it in case the protected LAN swells. Manage the protected LAN as usual. See PacketGuide's Control Strategy Overview for help

Scaling Considerations

The presence or absence of a firewall does not impact the choice of a model or ability to accommodate growth. If a unit is going on a core site's link, consult *Traffic Tree Options* to choose an appropriate traffic tree configuration for your needs. Then, no matter where a unit is being deployed, consult the Configuration Limits to select a model.

Core Site Internet Link Topology

A core-site Internet-link topology has a PacketShaper at the core site's link to VPN-connected branch offices, dial-up VPN users, partner extranets, and/or simply the web.

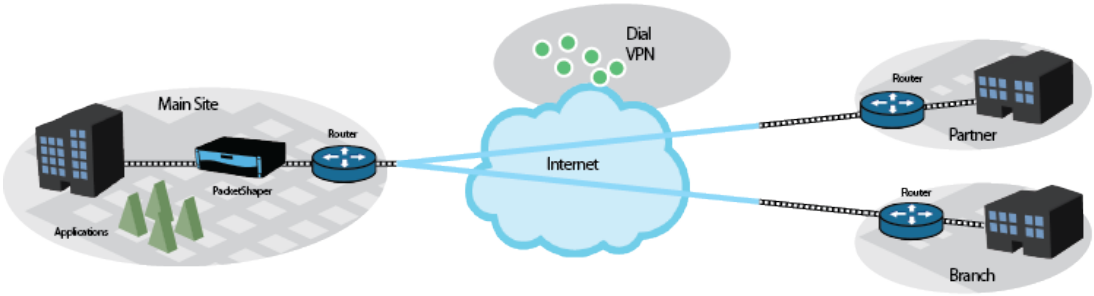

Environmental Requirements

There are no special network special requirements. Applications for VPN-connected branch offices must be hosted at core site servers.

Reasons to Choose

Customers choose this topology when their core site's Internet link supports both critical and casual traffic, or just to find out what traffic is on their Internet link and how it behaves. Deployed in this topology, PacketShaper can protect VPN traffic, protect dial-up users, contain unsanctioned web traffic such as music downloads, pace multimedia for optimum performance, reveal top web destinations, and more.

Installation Considerations

If you don't use VPN, installation is the same as the procedures described in the *PacketShaper Quick Start Guide*. As usual, you plug the Internet router's LAN port to the appliance's OUTSIDE port and plug the LAN to the appliance's INSIDE port. Then follow Guided Setup. If you use VPN, consult the *Virtual Private Network Topology* section that appears later in this document.

Capabilities

- PacketShaper's web classification features are particularly handy for this Internet link topology, as web traffic can vary in urgency, sensitivity to latency, and performance requirements. Classification based on travel direction, server location, mime type (XML, MPEG, and others), URLs, and other criteria help distinguish between casual browsing, business applications, customers' online activities, and so on.

- A PacketShaper offers all its monitoring and analysis features for all traffic that passes through the unit.

  For example, you can see how much of your Internet link goes to VPN branch offices and how much goes to casual browsing.

- When shaping is enabled, a PacketShaper offers all its control features on all traffic that passes through the unit.

  For example, you can allocate a third of your link to the branch VPN traffic or to your partners' extranet. Give each VPN user an appropriate minimum and maximum. Allocate a steady rate to multimedia streams. Block unsanctioned traffic such as music downloads or web sites with questionable content.

Limitations

A PacketShaper at only the core site does not aid performance at other branch offices with their own Internet links. It also does not aid performance for WAN traffic that uses a different link and network.

Configuration Suggestions

If you do not manage traffic to multiple branch offices on your Internet link, use a traffic tree that is based upon network applications, rather than branch locations or subnet. Automatic traffic discovery and classification are a good basis for the traffic classes, but you can also add or tailor classes based on other criteria. It's a simple tree to create and use, managing traffic according to type of traffic (such as application) rather than according to destination or location. Suggestions include:

- Customize an application-based tree with some of the web classification criteria to differentiate customers versus partners versus employees versus any other group you want to protect or control separately.

- Create a folder for streaming media and move all streaming protocols underneath it. Assign rate policies with guaranteed, per-session bandwidth. Consult PacketGuide's Recommendations for instructions.

- Create a folder called Unsanctioned or Undesirable or any other name that suits. Move any traffic classes that you would like to block (or at least contain to small amounts of bandwidth) to the folder. Then assign a policy to limit or block the traffic. Consult PacketGuide's Recommendations for instructions.

- If you use VPN, create a folder class to handle all VPN traffic types. Then assign a carefully sized partition to protect your VPN traffic. Use a dynamic subpartition on a VPN class if you want to provision bandwidth for each dial-up user.

- If you do manage multiple branch offices' traffic on your core site's Internet link, consider using instead the VPN topology.

Scaling Considerations

The PacketShaper model for a core site's Internet link needs to accommodate the link's capacity, anticipated maximum number of concurrent flows, and anticipated number of users. Although most appliances deployed in this topology use a traffic tree that is based primarily upon the applications running on the network, if you choose to create a traffic tree based upon branch office or subnet locations, the appliance needs to accommodate the number of traffic classes and matching rules required for the tree. Consult the Configuration Limits to check which model suits your tree's needs.

Core Site WAN Link Topology

A core-site WAN-link topology has a PacketShaper only at core sites.

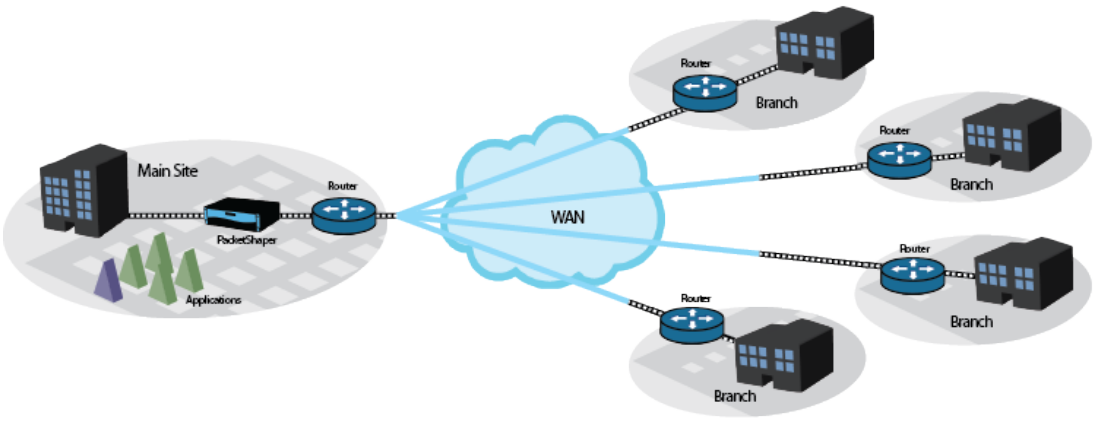

Environmental Requirements

To monitor and control performance most effectively from a centralized location, the WAN should be a pure hub-and-spoke design. All branch traffic must go through the core-site PacketShaper, and there should be no traffic between branch locations, which in this topology can reduce the unit's ability to calculate utilization for each link. Sometimes, the only way to determine that this is your exact network topology is to deploy a PacketShaper with the monitoring module at an branch office to check for traffic that doesn't head to or come from the core site. For example, many supposedly hub-and-spoke IP/Frame Relay hybrid networks are actually meshed networks where routing occurs inside a service provider's network.

Pure hub-and-spoke design is not necessary if you plan to manage only the core site's WAN link and not to manage branch offices' WAN connections. The more of the branch offices' traffic your PacketShaper can see, the more effective it can be. Three to one hundred branch locations are typical for this topology.

Application and Intranet servers must be at the core site's data center. Each branch location must access the Internet through:

- The core site, or
- A totally separate Internet connection that doesn't use the same last-mile connection to the WAN

Reasons to Choose

Customers with networks that fit the environmental requirements choose this topology to:

- Get visibility into network and application behavior before making further changes

- Try the technology before making a larger investment
- Avoid or phase the costs that come with additional branch units
- Deploy units gradually, as they are needed

Installation Considerations

Most core site WAN link installations require the same procedures as described in the *PacketShaper Quick Start Guide*. You plug the router's LAN port to the OUTSIDE port on the PacketShaper, and plug the LAN to the PacketShaper's INSIDE port. Then you follow Guided Setup.

Traffic Tree Options

A traffic tree is the construct that organizes and categorizes your traffic. Different types of traffic trees organize traffic differently – for example, by location or by application or both. Your choice of a style of traffic tree dictates:

- How PacketShaper will analyze and control your traffic
- Which features are accessible to you
- The configuration tasks you'll need to do
- How many branch offices and/or applications your model will support
- How easy it is to add managed branch offices or applications to your network

Consult the descriptions of traffic tree options and their implications in *Traffic Tree Options*. You are not limited to these trees, but they can provide insight into which options might be most appropriate for your own environment, and you can use them in initial strategies.

Topologies with Multiple LANs

This topology applies to any environment with multiple LANs connected to one or more routers attached to WAN or Internet links. Up to three LANs and three routers can be directly connected to a single PacketShaper.

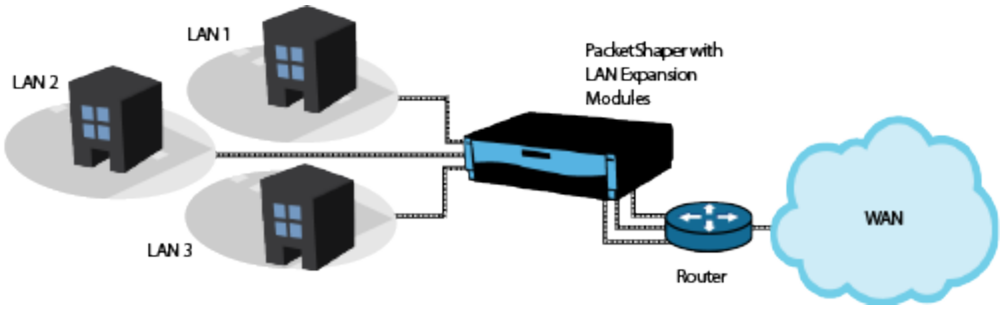

All traffic using any portion of access-link bandwidth should be managed together with a cohesive strategy. It is not effective to enforce policy-based bandwidth allocation on one portion of a router's traffic, if uncontrolled traffic from an additional LAN can swamp the router's capacity and undermine any potential performance gains. Symantec's network interface cards (NICs) extend PacketShaper benefits to organizations with multiple LANs without the need for extra switches or appliances to aggregate traffic.

Installation Considerations

Symantec offers copper and fiber-optic NICs for the PacketShaper S-series. Select the appropriate NIC for your model and link type.

- PS-S200 NIC Support
- "PS-S400 NIC Support " on page 943
- "PS-S500 NIC Support " on page 947

See the *Maintenance and Upgrade WebGuide* for installation instructions.

Capabilities

- A PacketShaper offers all its monitoring, analysis, and control features for all traffic on all NIC ports.
- All NICs include bypass . The bypass allows the PacketShaper to continue passing traffic between the two ports if the appliance fails or is powered off. See "Bypass Mode Overview" on page 69 for details.
- PacketShaper supports port link aggregation technology. NICs can be connected to an EtherChannel (trunks carrying multiple VLANs) but be aware not to overload the capacity of the PacketShaper.

Topology Examples with Multiple NICs

Each of the example topologies uses the following legend:

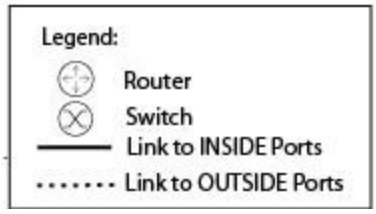

Example 1: In this topology, four LANs are attached to a single switch which is connected to the INSIDE ports on the PacketShaper NICs. All OUTSIDE ports are connected to a single router.

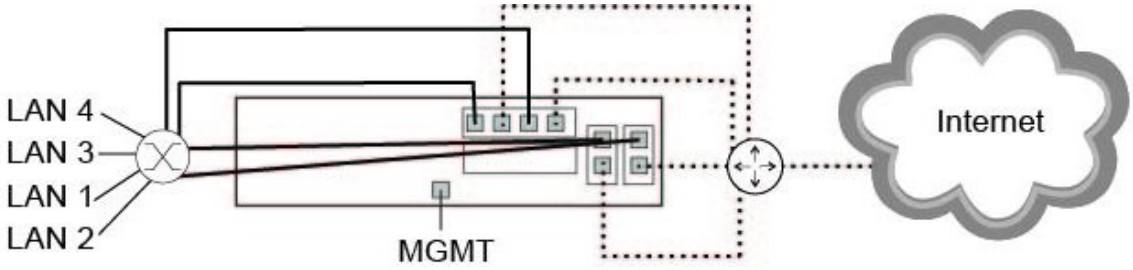

Example 2: In this example, the four LANs switches are connected to the INSIDE ports on PacketShaper NICs. Each OUTSIDE port is connected to a corresponding WAN router.

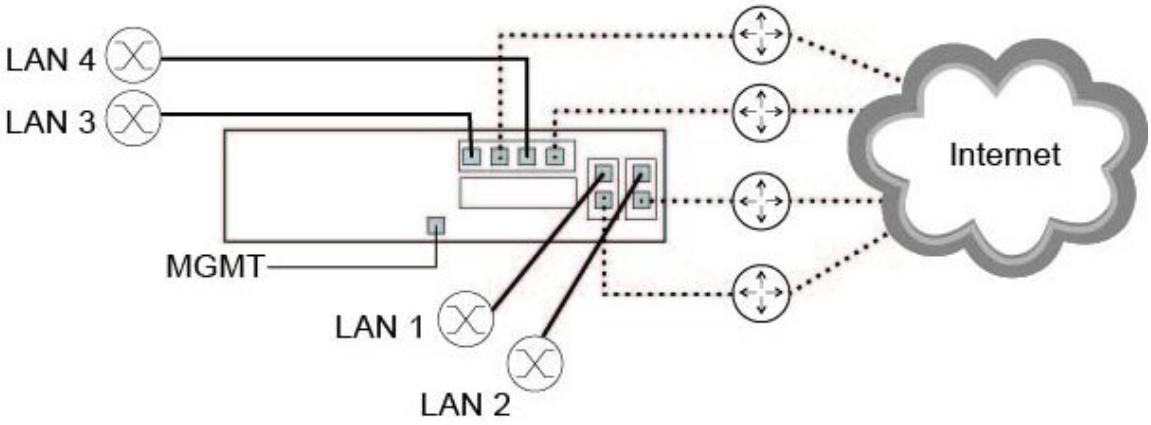

Limitations

PS-S200 can have a dual- or quad-port NIC installed in Slot 3. If a quad-port NIC is used, the PS-S200 can accommodate up to three LANs.

PS-S400 can have dual- or quad-port NICs installed in Slots 3 and 4. If these slots are populated with quad-port NICs, the PS-S400 can accommodate up to five LANs.

PS-S500 supports installation of NICs in Slots 1-4. If Slots 3 and 4 are each populated with quad-port cards, up to six LANs can be accommodated on a single PS-S500.

Configuration Suggestions

For the topology shown in example 1, you can separate statistics and control strategies for each LAN by defining subnet- or IP address-based traffic classes for those LANs and then configuring a static partition for each class. For the topology shown in example 2, you can classify traffic by NIC (for example, Slot1, Slot2, Slot3) to segregate LAN traffic in the tree. For either topology, the remainder of your management strategy is not impacted by the presence or absence of multiple LANs.

Scaling Considerations

Choose a model based on total WAN-link capacity, not the needs of one particular LAN. Be sure to verify if the model that suits your capacity needs supports your desired NIC. Consult the Configuration Limits.

High Availability Topologies

In Symantec's high availability solution for PacketShaper S-Series, appliances can be installed in redundant network paths to provide link redundancy in case a component (PacketShaper, router, or switch) fails. In addition, the PacketShaper high availability solution can be applied in a network environment where asymmetric traffic tends to occur. All PacketShaper units in a configured high availability cluster are active and manage traffic on their own links, but they are simultaneously and continuously monitoring the health and status of its partners to make sure a unit has not gone down.

The example diagram below shows two PacketShaper S-Series appliances (different models) in a high availability group or *cluster*. Note that each PacketShaper manages its own link and has network routes to one another, but doesn't need to be directly connected by a cable to other partners.

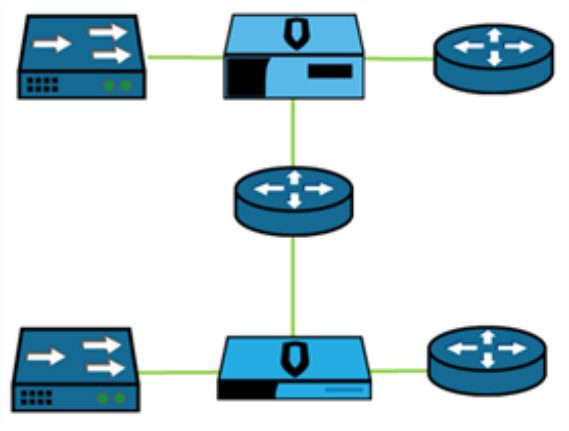

Installation Considerations

To make the PacketShaper function in these types of topologies, you connect each unit to each redundant network path, and make sure you can still access each unit if a router or switch goes down with a *redundant management path* from the unit to the switch on the parallel path. The topology drawings below detail connections in several topologies.

Topology Examples

The example topologies show two PacketShaper partners, but a cluster may have up to five partners. The topologies use the following legend:

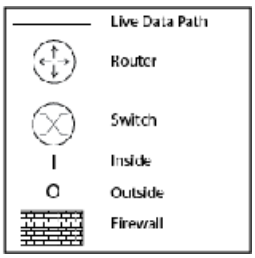

Basic Failover Topology

The basic topology uses two switches and two routers. If a device fails in one path, the devices in the redundant path take over to handle all the traffic.

Before PacketShapers:

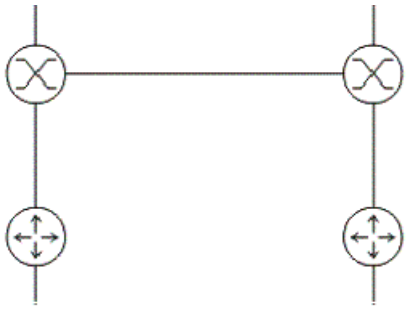

With PacketShapers:

This example shows two PacketShaper partners in a cluster; one is a PS-S500 and the other is a PS-S400. Note that partners can be any combination of PacketShaper S-Series appliances, and you can have up to five partners in a cluster.

The SYS MGMT port is connected to a management network to provide an alternate way to manage the unit.

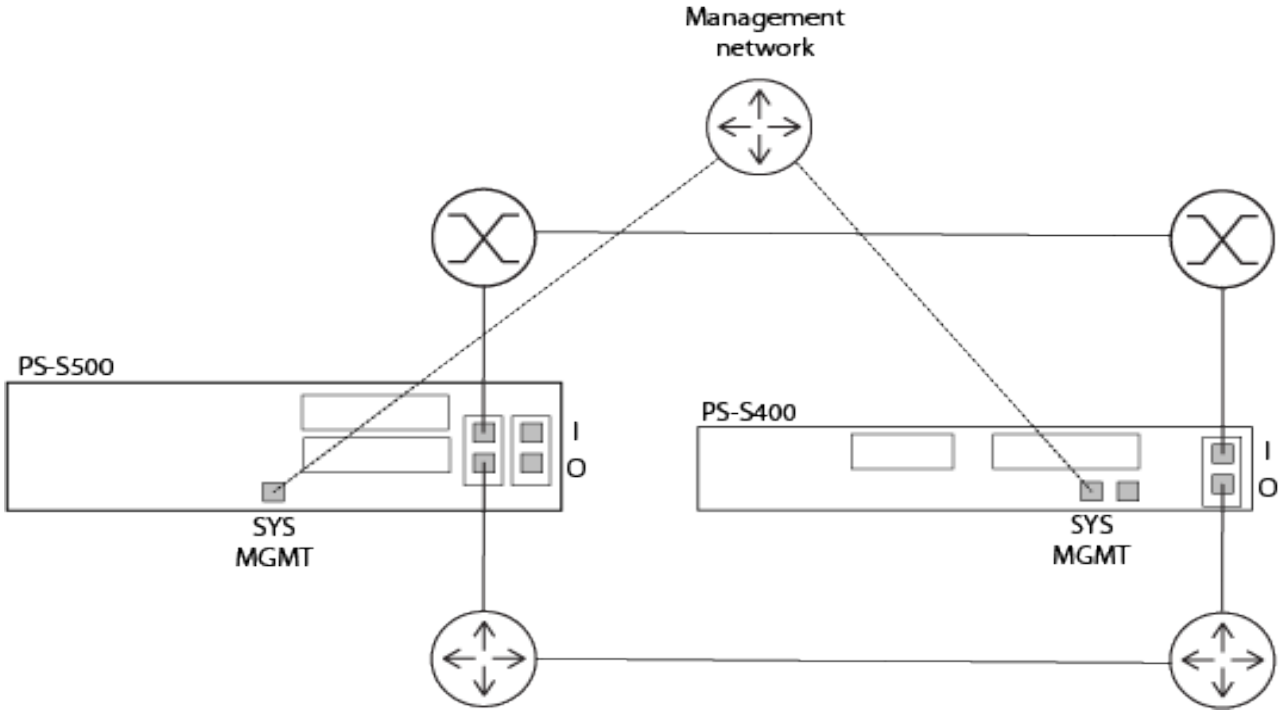

Load Balancing Topology

The load balancing topology shown below controls two links on the same VLAN or subnet. It uses two switches and two routers, with each router connected to both switches to allow load balancing between the two paths. If a device fails, the devices in the redundant path take over to handle the traffic.

Before PacketShapers:

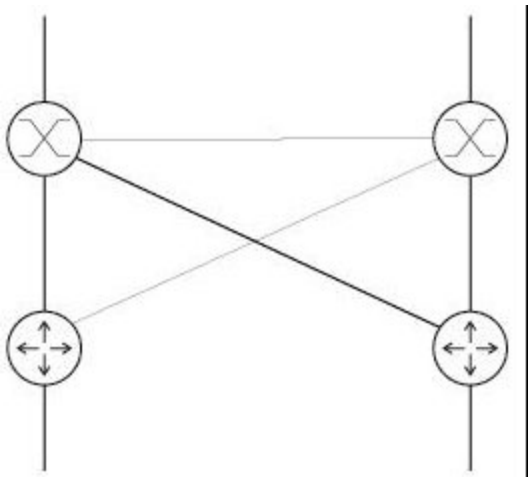

With PacketShapers:

This example shows two PacketShaper partners in a cluster; one is a PS-S500 and the other is a PS-S400. Note that partners can be any combination of PacketShaper S-Series appliances, and you can have up to five partners in a cluster.

This topology requires a NIC on each PacketShaper to connect to the redundant data path. On a PS-S500, you can use the NIC in Slot 2. On a PS-S400 or PS-S200, you will need to purchase an option card. To provide an alternate way to manage the unit, use the SYS MGMT port for connecting to a management network.

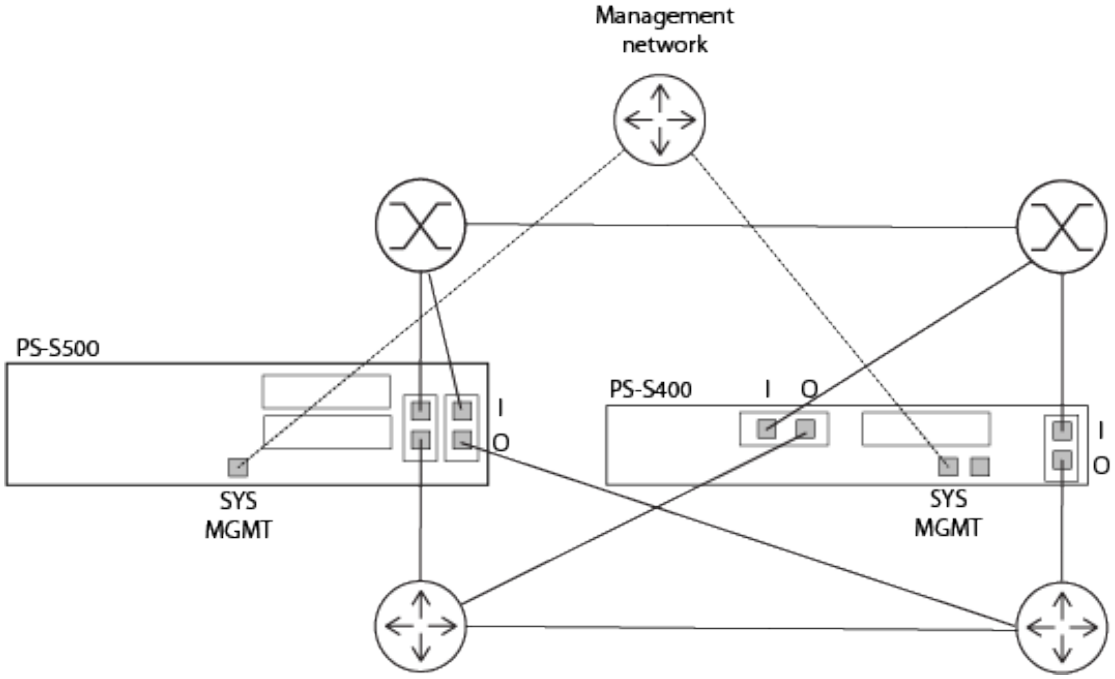

Load Balancing with Multiple VLANs

The load balancing topology shown below controls two links on multiple VLANs or subnets. It uses four routers, with each router connected to the other router on the alternate path. If a device fails, the devices in the redundant path take over to route and handle the traffic.

Before PacketShapers:

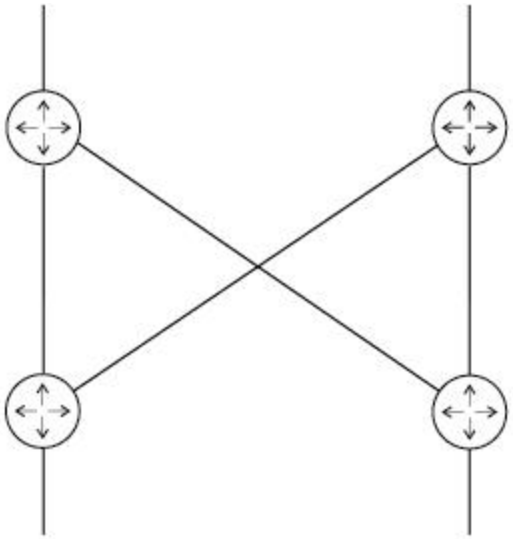

With PacketShapers:

This topology requires a NIC on each PacketShaper to connect to the redundant data path. On a PS-S500, you can use the NIC in Slot 2. On a PS-S400 or PS-S200, you will need to purchase an option card.

To provide an alternate way to manage the unit, you can use the SYS MGMT port for connecting to a management network.

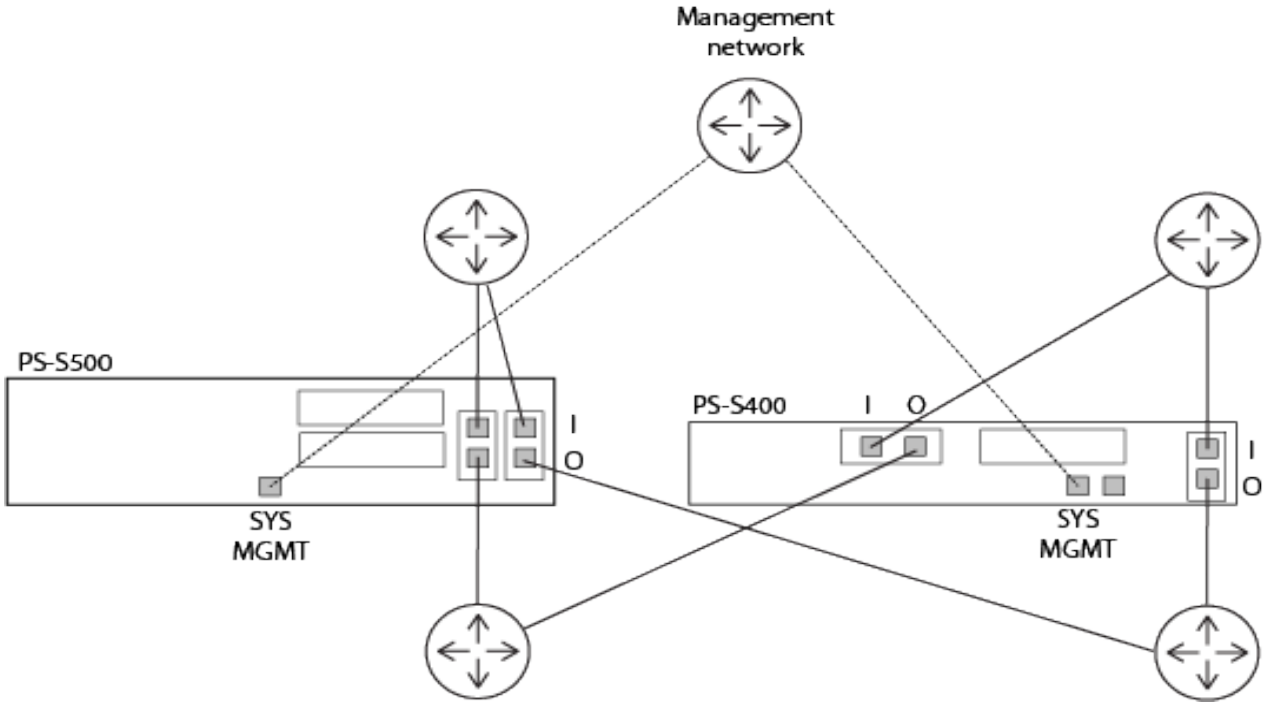
Management
network
PS-S500
I
O
SYS
MGMT
PS-S400
I
O
I
O
SYS
MGMT

Topologies with Firewalls

The firewall shown in this topology can actually be any generic device that sits in this position.

Before PacketShapers:

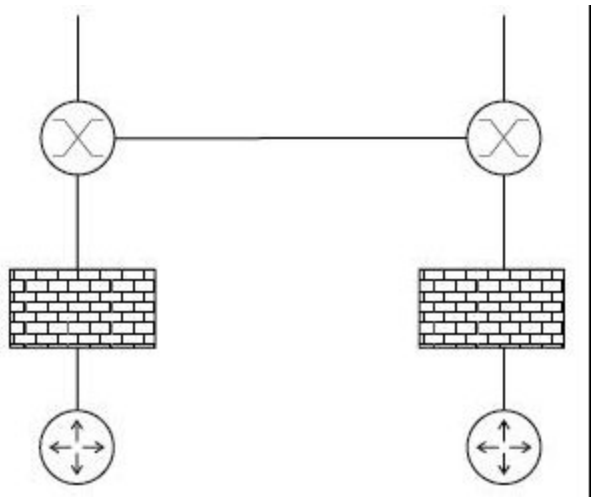

With PacketShapers:

The OUTSIDE ports on Slot 1 connect to firewalls. The SYS MGMT ports are connected to a management network to provide an alternate way to manage the units.

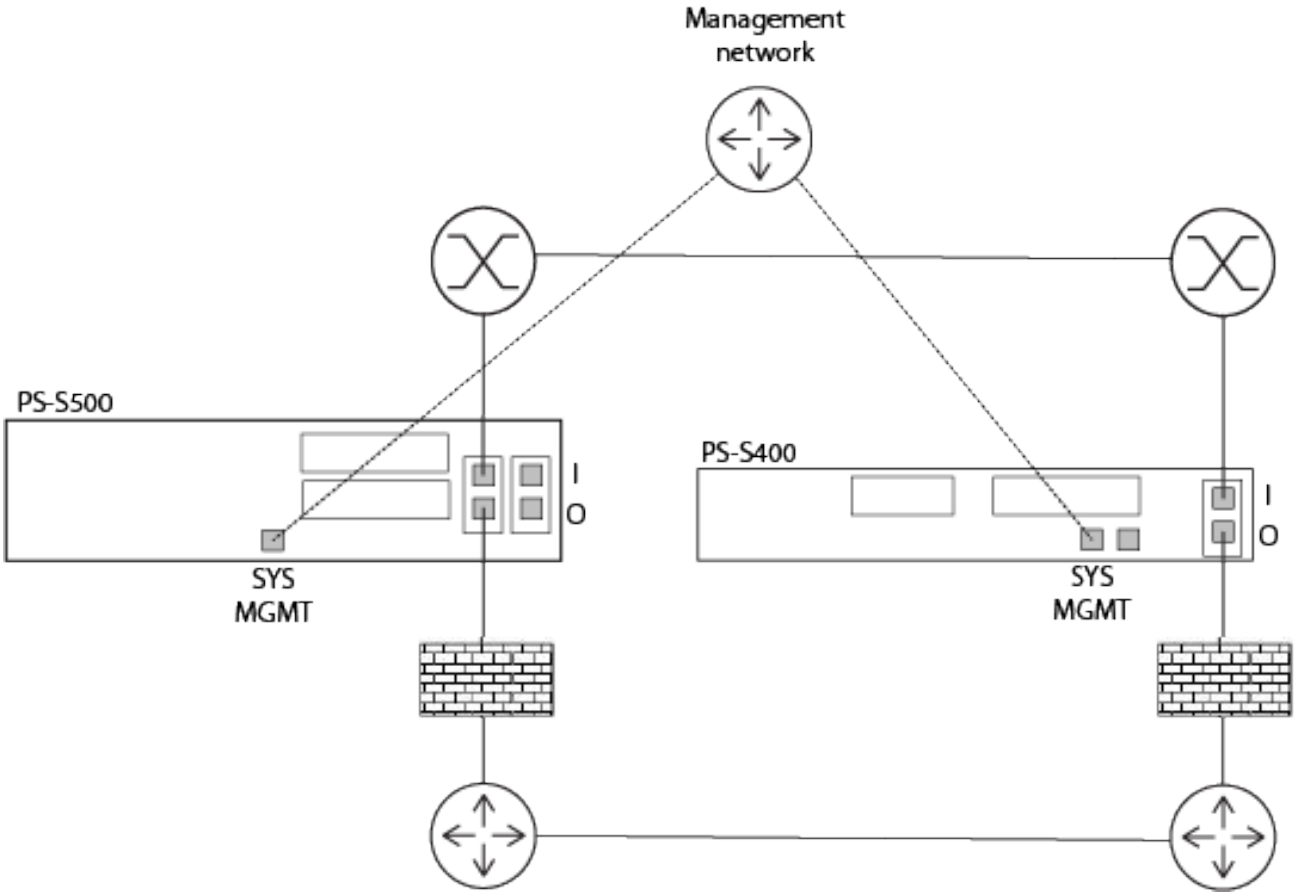
Management
network
PS-S500
I
O
SYS
MGMT
PS-S400
I
O
SYS
MGMT

Capabilities

- All PacketShaper units in a configured high availability cluster are active and manage traffic on their own links, but they are simultaneously and continuously monitoring the health and status of its partners to make sure a unit has not gone down.
- PacketShaper offers all its monitoring and analysis features, including classification, utilization analysis, service-level management, efficiency analysis, thresholding and events, packet capture, diagnostic graphs, and reports.
- With shaping enabled, PacketShaper is able to protect, contain, pace, and provision bandwidth, and provide TCP Rate Control, mark packets, and suppress DoS attacks.
- Asymmetric flows are supported.

Limitations

- When standby partners are down, the active and standby units share available bandwidth. Ten percent of available bandwidth is allocated to each standby device that is down. Thus, if one standby partner is down, the active system uses 90 percent of the available bandwidth and the failed device gets 10 percent. For each standby partner that is down, PacketShaper creates a partition with a size equal to 10 percent of the WAN link size.

Configuration Suggestions

As usual, protect critical applications, pace less urgent applications, and contain unsanctioned applications.

Scaling Considerations

Keep in mind that when a PacketShaper is taking over for a failed link, it is processing the traffic going to both routers and both WAN links. Therefore, assuming both links are active, each of the units needs to be a model that can process the amount of throughput in *both* routers. For example, if you have two 45 Mbps links, each with its own router and each active, then each PacketShaper needs to be able to support 90 Mbps.

Traffic Tree Options

The way PacketShaper organizes traffic dictates how it can analyze and control that traffic. Therefore, your traffic tree configuration is a crucial choice in determining the features you have available. For example, if your tree doesn't distinguish SAP from Oracle, then you can't measure bandwidth utilization for each. Or, if your tree doesn't distinguish traffic to your Paris office from traffic to your Oslo office, then you won't be able to compare traffic volumes or give prescribed amounts of bandwidth to each branch office.

Traffic trees are very flexible. You can customize your own trees' classes using many types of criteria including application, location, protocol, host lists, addresses, bit markings such as ToS, MAC address, and many other strategies. PacketShaper also allows you to classify network traffic by physical port or NIC, where each NIC may have its own subset of defined traffic classes. We suggest a few options of tree types with their implications here. You are not limited to these trees, but you can use them in initial strategies.

To help determine which traffic trees might suit your own circumstances, have the answers to these questions ready before reading the tree descriptions.

- How many branch offices do you manage?
- How many applications (or other subsets of traffic, such as all traffic to one particular Oracle database) do you want to manage separately? For each remote office?
- Can you fit your applications into 3 to 5 categories of management strategies? Would it be okay if you had usage statistics and graphs only for each category (Critical, for example) instead of for each application (SAP, for example)?
- Do you want separate usage statistics and graphs for each application? For each application at each branch office?
- Do you want to control each application for each branch office distinctly? Or can one policy/partition strategy apply to each application, no matter the application's branch location? For example, does Citrix need a different policy at one branch office than at another? Perhaps because the different branch office's link capacities vary?

1. Application-Based Traffic Tree

This tree is for a WAN or Internet link at any core or branch site where you want to vary traffic management strategies according to type of traffic (such as application) rather than according to destination or location.

An application-based tree is most common at branch offices and core sites' Internet links. Automatic traffic discovery and classification are a good basis for the traffic classes, but you can also add or tailor classes based on other criteria. It's a simple tree to create and use.

Capabilities

A PacketShaper offers all its monitoring and analysis features for any traffic class. For example, you can compare bandwidth usage for SAP and Oracle, but not for Sao Paulo and Bogota.
When shaping is enabled, a PacketShaper offers all its control features, such as partitions and policies. For example, you can allocate 512 Kbps to all of SAP and a steady 22 Kbps to each VoIP session.

Limitations

No insight into or control over per-location behavior.

Scaling Considerations

Choose a product model based on your link's capacity, anticipated maximum number of concurrent flows, and/or number of users. This configuration rarely bumps into limits on traffic classes or matching rules. Consult the Configuration Limits to check which model suits your needs. There is less growth-associated overhead than for location-based trees. Automatic traffic discovery detects many new applications and creates classes for you.

Other Variations

- **Category-based tree**: monitors and controls traffic based on URL categories.
- **Model tree**: a pre-configured tree that organizes network traffic into folders of typical genres, such as VoIP, risky websites, business-critical applications and data, and recreational websites and applications.

2. Simple Location-Based Tree with Locations Only

This tree is for a core site's WAN or Internet link with traffic that goes to multiple branch offices. It categorizes first by travel direction and then by location. It's appropriate for occasions when your primary concern is to provision bandwidth, and you don't

particularly care how it's used or how applications perform. Additionally, it can be appropriate for a core site in topologies where other PacketShaper appliances manage applications at each location, and the core site's appliance just manages the amount of traffic going out to each branch office.

Capabilities

- A PacketShaper offers all its monitoring and analysis features — ***but***, on a global and per-location basis only. For example, you can compare bandwidth usage in Sao Paulo with usage in Bogota, but you can *not* see how much of Bogota's usage is music downloads or SAP.

- When shaping is enabled, a PacketShaper offers all its control features — ***but***, on a global and per-location basis only. For example, you can allocate 512 Kbps to Bogota, but you can *not* block music downloads or protect SAP.

Limitations

No insight into or control over per-application behavior. No visibility into which applications travel your network.

Scaling Considerations

This management option scales well, as it accommodates many locations easily. Choose a product model based on your link's capacity, anticipated maximum number of concurrent flows, and number of traffic classes to accommodate your number of locations (times 2, to have each location under both Inbound and Outbound). Consult the Configuration Limits to check which model suits your needs.

There is little overhead involved in adding branch offices, as you need only to create an additional class under both the Inbound and Outbound classes.

Configuration Suggestions

The most common criteria to define location-based classes are subnets or host lists. You can create initial location classes one-by-one by hand, or, if there are many classes and it gets to be a tedious task, you can automate the process with a command file and a data entry form.

At a core site, give each branch office a static partition with a size equal to the branch office 's WAN-link capacity. If you are concerned with each user's bandwidth, define a dynamic, per-user partition under each branch office's static partition.

3. Location-Based Tree with Per-Location Applications

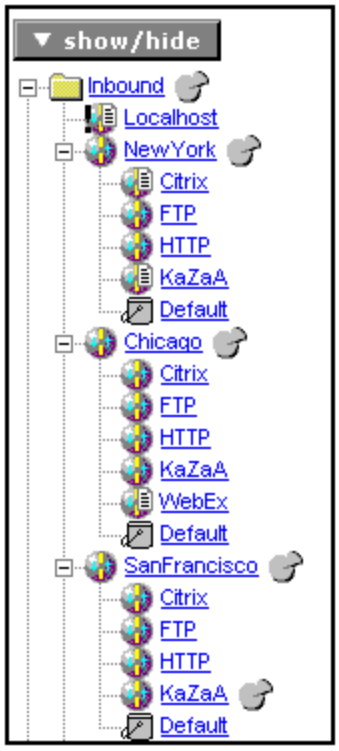

This tree is for a core site's WAN or Internet link with traffic that goes to multiple branch offices. A location-based tree with per-location applications gives you the most insight and the most control over your applications. But it imposes scaling considerations and takes more time to configure. It categorizes first by travel direction, then by location, and finally by application.

Capabilities

- A PacketShaper offer all its monitoring and analysis features.

  This is an appropriate configuration choice when you want the most detailed level of analysis. For example, if you want to know how much New York traffic you have, *and* you want to know how much SAP traffic you have, *and* you want to know how much SAP traffic you have in New York – then this traffic tree is your best bet.

- When shaping is enabled, a PacketShaper offers all its control features.

  This is an appropriate configuration choice when you need to control the traffic to each location on a per-application basis. For example, suppose the Sao Paulo branch office has more capacity than Bogota has. You don't want to limit Sao Paulo's FTP to match the level that Bogota can handle. But you also do not want to burden Bogota with FTP just because Sao Paul can handle it. With this configuration, Sao Paulo and Bogota can each have different FTP limits.

Limitations

You must be prepared to manage the details of each location separately. You must assign appropriate policies for each application you wish to manage at each location.

Scaling Considerations

This style of traffic tree tends to use up many traffic classes: (number of applications+1(for default)) * (number of branch offices+1) * 2(for inbound and outbound). Check to make sure your total fits within your model's capacity in the Configuration Limits.

For example, if 15 branch offices each have 20 applications, 16 * 21 * 2 = 672 classes. You would need a PacketShaper that can accommodate this tree, even if your WAN speed doesn't require that model's capacity.

If you need or will shortly need more traffic classes than are available, you have outgrown a core-site-only topology. Instead, choose a distributed topology with a simple location-based tree type at your core site.

If you add a location, you add just the one class under Inbound and Outbound. Then you can use automatic discovery features (or you can copy other branch offices' classes). If you add an application, automatic discovery will probably find it. Or, you can copy a manually created class. You must assign appropriate policies and partitions to each of an application's many classes.

Configuration Suggestions

At a core site, give each branch office a static partition with a size equal to the branch office's WAN-link capacity. Then, under each branch location, assign partitions and policies to each application that needs protection or that undermines other applications' performance. Use minimum and maximum sizes that fit appropriately under the associated branch office's partition. For example, if New York has a 2 Mbps partition, then you could give SAP 500 Kbps, but not 3 Mbps. Consult Policy and Partition Guidelines for specific suggestions on specific applications and Control Strategy Overview for more high-level advice.

4. Location-Based Tree with Application Groups

This tree is for a core site's WAN or Internet link with traffic that goes to multiple branch offices. A location-based tree with application groups categorizes first by travel direction, then by location, and finally by group of application. The groups can be something similar to: Urgent, Average, NonUrgent, Prohibited, and Default. Or perhaps you prefer groups such as Protect, Pace, Contain, and Prohibit. Any grouping strategy in which each group is managed with one strategy works.
This option offers some insight and control over your applications – more than the locations-only option, but less than the per-location-applications option. Compared to the per-location-applications option, it takes the same or more effort to configure, uses fewer traffic classes and more matching rules, and is easier to view and to understand.

Capabilities

- A PacketShaper offers all its monitoring and analysis features – ***but***, on a global, per-location, and per-group basis only.

  For example, you can graph your total usage for New York. But you can*not* know how much bandwidth SAP in New York uses.

- When shaping is enabled, a PacketShaper offer control features – ***but***, on a global, per-location, or per-group basis only.

  For example, you can protect bandwidth for your urgent applications to New York, reserving perhaps 50 percent of New York's 512 Kbps capacity. But you can *not* reserve bandwidth only for SAP going to New York. You can indeed block music downloads to New York, even at the same time as allowing them to flow to and from your Los Angeles office.

Limitations

No per-application insight or control. Defining groups with all required matching rules can be a long process.

Scaling Considerations

This style of traffic tree tends to use up many matching rules, the construct that defines which types of traffic match which traffic classes. Each application can have several types of associated traffic, requiring several matching rules. In addition, each type of TCP traffic can require two matching rules, one for an inside server and one for an outside server (if both are possible).

For example, if one of your Urgent traffic classes matches SAP and Oracle traffic, it would need matching rules for the following types of traffic:

- Oracle-netv1 Oracle SQL*Net V1-based traffic (v6, Oracle7)
- Oracle-netv2 Oracle SQL*Net V2/Net8-based traffic (Oracle7, 8, 8i or 9i)

- Oracle-SSL Oracle database over SSL
- Oracle-JVM-SSL Oracle JVM traffic (IIOP) over SSL
- OracleClient Oracle Java client (WebForms)
- OracleEM1 Oracle Enterprise Manager-1
- OracleEM2 Oracle Enterprise Manager-2
- SAP Systemanalyse und Programmentwicklung

Or, if an urgent traffic class included Telnet and TN3270, it would need matching rules for the following types of traffic:

- Telnet-Clear Telnet terminal traffic
- Telnet-Secure Secure Telnet terminal traffic (SSL)
- TN3270 IBM 3270 terminal traffic over telnet

You can estimate your matching-rule needs with: (the sum of all matching rules for one branch office) * (number of branch offices) * 2 (for inbound and outbound). Check to make sure your total fits within your model's capacity iin the Configuration Limits.

For example, suppose 30 branch offices each have four groups. The parent class (for the branch office) needs one IP address-based matching rule. Assume the Urgent traffic class requires 15 matching rules, Average requires 10, NonUrgent also needs 10, Prohibited needs just three, and Default has just one matching rule. ([1+15+10+10+3+1] * 30 * 2) = 2400. You would need a PacketShaper with at least 2500 matching rules to accommodate this tree. If many of your matching rules require OSI Layer-7 classification, consider a higher-capacity unit because this type of classification requires more processing power.

If you need or will shortly need more matching rules or traffic classes than are available, you have outgrown a core-site-only topology. Instead, choose a distributed topology with a simple location-based traffic tree at the core site.

To add locations, you need only to add the class under both Inbound and Outbound. The group classes need only be created once. A bulk-copy operation eases the overhead of copying them to many branch classes. If you add an additional application to one of the groups, you must touch many group classes to add the appropriate matching rule(s).

Configuration Suggestions

At a branch location, give each branch office a static partition with a size equal to the branch's WAN-link capacity. Then, under each branch location, assign a partition and policy to each group. For example, give Urgent a partition of 25 to 75 percent of your branch office's capacity, depending on the urgent applications' bandwidth needs. Assign a policy to block traffic in your Prohibited class. Give your Average class steady access to bandwidth but prevent it from swelling to undermine others' performance. And so on. Consult Policy and Partition Guidelines for assistance.

5. Location-Based Tree with Global Applications

This tree is for a core site's WAN or Internet link with traffic that goes to multiple branch offices. A location-based traffic tree with global applications categorizes traffic first by travel direction, then by location. That's all. The metrics, graphs, and reports are the same as those available in the simple location-only tree. The big difference is that you can define separate bandwidth-

allocation policies for each application. But all traffic from one application is managed with one strategy, independent of branch location.

This option is most common when all branch offices have the same link size and similar business priorities because each application is managed the same way for all locations. This traffic tree is the most difficult to understand conceptually, but the easiest to configure and needs the least amount of traffic classes, matching rules, and configuration time. It provides a very scalable and easily managed core-site tree for traffic shaping but does limit visibility.

This type of tree requires that shaping be enabled because of its use of policies.

Capabilities

- A PacketShaper offers all its monitoring and analysis features — ***but***, on a global and per-location basis only.

  For example, you can compare bandwidth usage in Sao Paulo with usage in Bogota, but you can *not* see how much of Bogota's usage is Citrix. You can *not* see which are the top web sites.

- When shaping is enabled, a PacketShaper offers all control features on a global and per-location basis, and it offers policies on a per-application basis.

  For example, you can allocate 512 Kbps to Bogota, and you can assign TN3270 a priority of 7. But you can *not* assign TN3270 a priority 7 for Bogota and a priority 3 for Quito.

Limitations

No per-application insight. Must use inheritable policies, one of the product's more conceptually difficult constructs.

Scaling Considerations

This management option scales well, as it accommodates many locations and applications easily. Consult the Configuration Limits to check which model suits your WAN link speed and has the number of traffic classes to suit your needs: ([number of branch offices] + [number of applications]) * 2 (Inbound and Outbound).

For example, if you have 20 branch offices and 40 applications, (20 + 40) * 2 = 120 traffic classes, well within the means of all models. You would use link speed or another option to choose your model.

There is little overhead involved in adding branch offices. You need only to create an additional class under both the Inbound and Outbound classes. Automatic discovery will probably take care of adding applications. And if not, you need only to create a traffic class once for each application under both Inbound and Outbound, as the classes are not repeated under each branch location.

Inheritance

For a complete description of inheritance, see Inheritance Rules. A summary of how inheritance applies in this type of tree follows:

The PacketShaper appliance classifies each passing traffic flow in the class that matches the flow's branch location. The appliance adds the metrics for the traffic flow only to the matching branch class. It applies any partition assigned to the matching branch class. But when the appliance finds no policy assigned to the branch class, it continues searching until another matching class with a policy is found.

If the tree is configured correctly, PacketShaper finds an application class as the second match. The flow *inherits* the second class' policy. But the second class gets no performance metrics recorded other than an additional policy hit.

For example, consider the following tree fragment with three branch offices and three applications:

Inbound
London (Partition of 512 Kbps)
Paris (Partition of 512 Kbps)
Rome
SAP (Rate Policy 6)
WebBrowsing (Rate Policy 3)
TN3270 (Priority Policy 7)

When a 100 KB TN3270 flow heads to London, the London class gets those bytes added to its usage figures, the flow is paced to fit within London's partition, and the flow gets top access to the partition's capacity with its inherited priority 7 policy. But the TN3270 class' usage is zero because the London class was the initial classification match. The TN3270 class' policy hits does increase by one. You can generate insightful reports about the London class as a whole, but not about the TN3270 class.

Configuration Suggestions

At a core site, give each branch office a static partition with a size equal to the branch office's WAN-link capacity. Do not assign a policy, as traffic will inherit its policy from the application classes. Then, create inheritable classes (beneath the list of locations) for each application and assign appropriate policies. Consult Policy and Partition Guidelines for suggestions on specific applications.

Virtual Private Network Topology

A virtual private network (VPN) is a private data network that uses the public telecommunication infrastructure to extend a WAN from a core site to branch offices or dial-up users. Encryption and security procedures keep data private. A VPN enables companies to share distributed data and applications without the expense of dedicated leased lines, frame relay, or private lines.

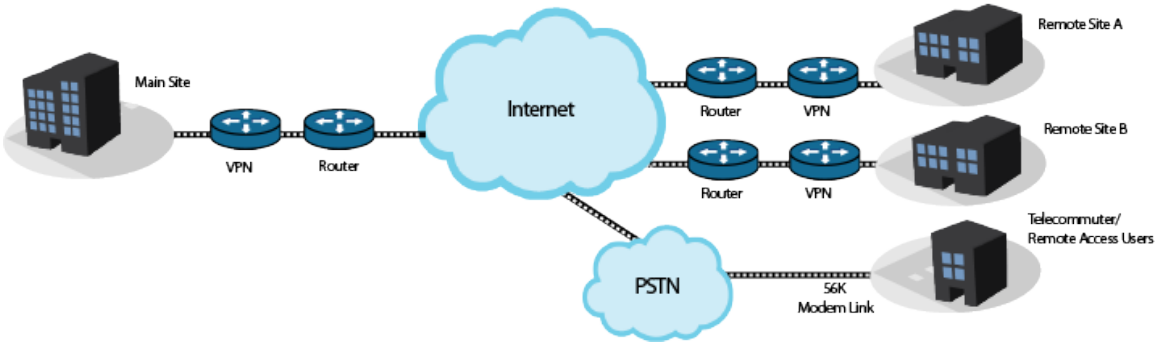

Although the Internet can serve as an economical WAN alternative, your enthusiasm for this solution can wane if you are faced with unpredictable performance for mission-critical applications traversing the VPN. To add to this problem, most monitoring tools aren't able to portray VPNs' encrypted traffic with any detail. A PacketShaper with shaping enabled assists both issues and PacketShaper with just monitoring (no shaping) can help in the latter.

Environmental Requirements

There are no special requirements.

Reasons to Choose

Customers choose this topology to protect the performance of applications on a VPN.

Installation Considerations

You can choose to install a PacketShaper on either side of your VPN gateway. Its position with respect to the VPN gateway dictates whether it will classify applications before or after encryption. The two deployment topologies for Internet links using VPN are shown below.

Example 1

Put the PacketShaper between the router and the VPN gateway if you want an easy way to spot all encrypted VPN traffic in order to protect it. In this example, you plug the Internet router's LAN port to the PacketShaper's OUTSIDE port and plug the VPN gateway to the PacketShaper's INSIDE port.

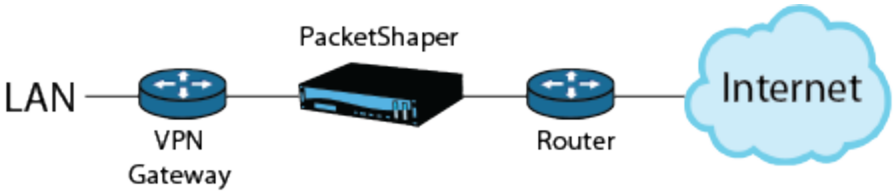

Example 2

Put the PacketShaper on the inside of the VPN gateway if you want it to differentiate between multiple encrypted applications. You plug in the VPN gateway's LAN port to the appliance's OUTSIDE port and plug the LAN to the appliance's INSIDE port.

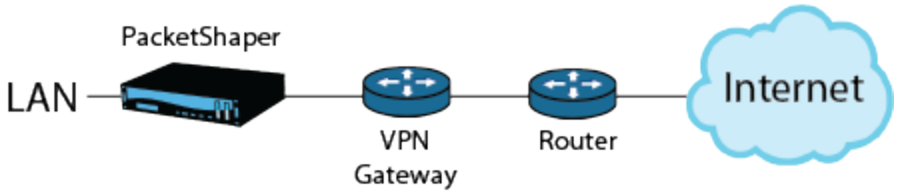

Capabilities

If you place the PacketShaper between the VPN gateway and the router, it sees both encrypted VPN traffic and normal web traffic or other applications. It classifies VPN traffic according to encrypted protocol and does not see which applications are encrypted. It automatically differentiates many types of secure traffic including IPSec (AH, ESP), PPTP, L2TP, DLS, DLA, SSL, SSH, RADIUS, GRE, IKE, SSL, ISAKMP, and SOCKS.

A PacketShaper can monitor and analyze for the traffic that passes through it. For example, you can see how much of your link goes to each VPN-connected branch office.

When shaping is enabled, PacketShaper offers additional control features. For example, you can allocate a set amount or percentage of your link to each branch office and a per-user amount of bandwidth for each dial-up user.

Configuration Considerations

If you installed the appliance so it can only see encrypted traffic, then use a simple location-based tree (consult Traffic Tree Options for details on this tree type). Assign a partition to each location matching the capacity of the branch office.

If you installed the appliance so it sees traffic before it is encrypted, create a location-based tree with per-location applications (consult Traffic Tree Options for details on this tree type). Then manage each application at each location separately with appropriate policies and partitions.

Scaling Considerations

The product model for a VPN needs to accommodate the link's capacity, anticipated maximum number of concurrent flows, and the number of traffic classes and matching rules required for the location-based tree. Consult the Configuration Limits to check

which model suits your tree's needs.

Non-Inline Deployments (Watch Mode)

PacketShaper's *watch mode* is used to monitor traffic non-inline, where the PacketShaper is not cabled into the core data path. This type of deployment is best suited for enterprises that have strict change control procedures or security restrictions on the introduction of inline elements into the network.

A PacketShaper appliance configured in watch mode passively monitors the traffic on the network and performs all the traffic classification and reporting tasks that the appliance typically does, providing insight into what applications are being used on the network and highlighting performance issues. A PacketShaper in watch mode cannot perform traffic shaping.

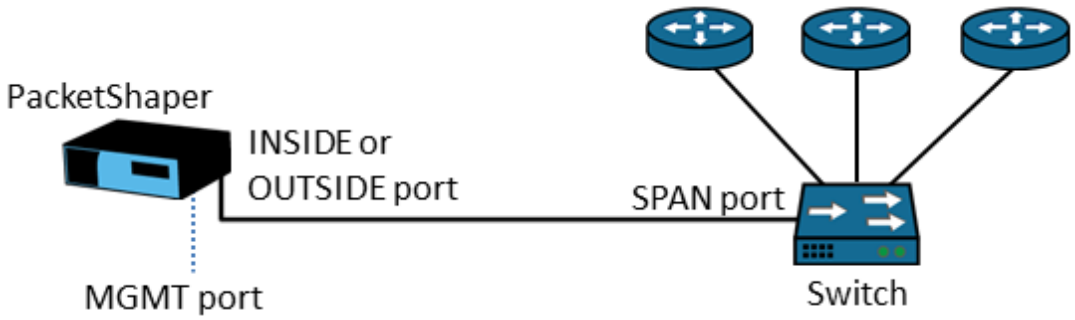

Reasons to Choose

A PacketShaper in watch mode can easily conduct network audits or managed application assessments while offering increased deployment flexibility in restricted or complex data center environments. The watch mode feature offers functionality similar to a network probe, but PacketShaper's extensive monitoring capabilities can classify and record measurement data by application (layer 7), not just by port.

Installation Considerations

A PacketShaper in watch mode works with any hub or switch; when connected to a switch, it should be connected to a SPAN port so that it has access to all the WAN traffic. If you want to monitor traffic sent by the switch to the router, you need to have the bidirectional traffic on the switch's router port copied over to the SPAN port in order for the PacketShaper to receive and monitor it.

Topology Examples

Example 1: Monitor One LAN

In this basic topology, a PS-S500 monitors network traffic via a connection between the INSIDE or OUTSIDE port on Slot 1

and a switch's SPAN port. The SYS MGMT port is used for PacketShaper access and management, and must be connected to the appropriate LAN segment for PacketShaper access. Traffic received on the management port is not classified or measured.

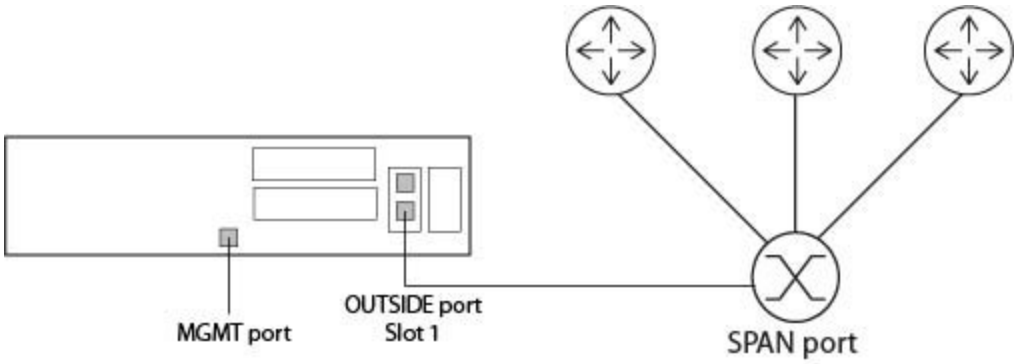

> **Note:** Although the example above shows the OUTSIDE port connected to the SPAN port, you can use either the INSIDE or OUTSIDE port.

On PS-S200 and PS-400 models, Slot 2 is the built-in NIC that would be used to monitor a single LAN segment.

Example 2: Monitor Multiple LANs

In this example, a PS-S500 monitors traffic on three different LAN segments. This is achieved by using NICs. In the diagram below, the OUTSIDE port on Slot 1 is connected to a switch or hub on one LAN and the OUTSIDE ports on NICs are connected to two other LANs. Although OUTSIDE ports are shown in this example, you can connect to either the INSIDE or OUTSIDE port. You can use the SYS MGMT port to manage the unit from another LAN.

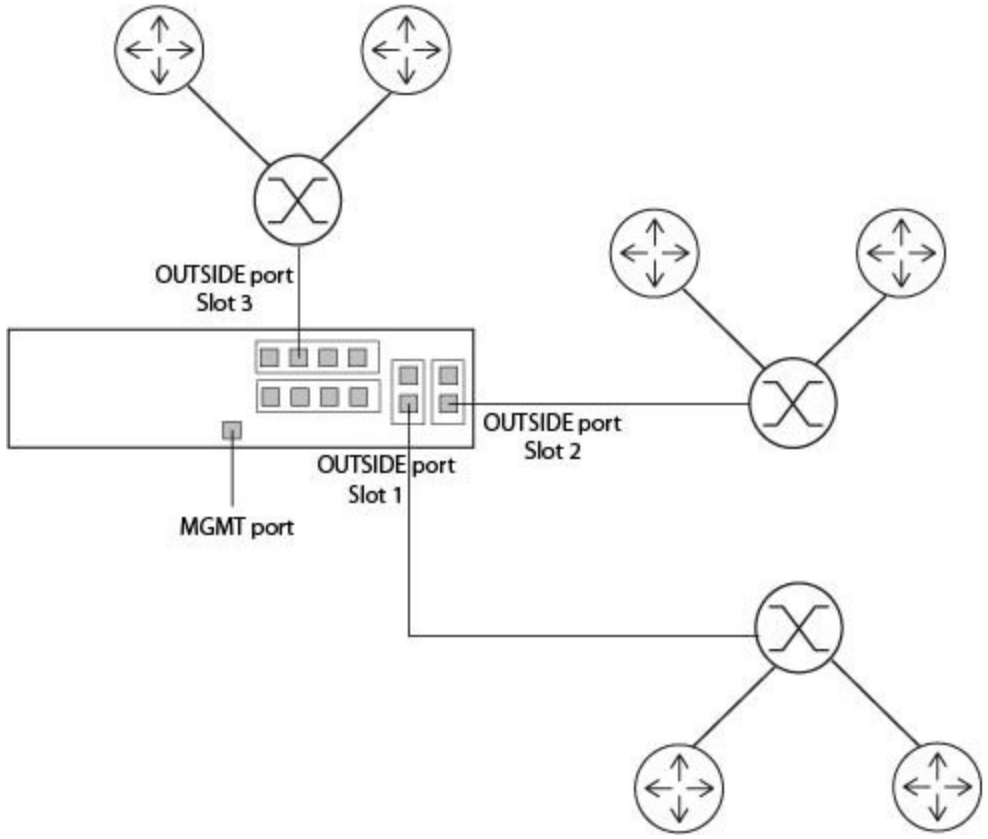

PS-S200 can have a dual- or quad-port NIC installed in Slot 3. If a quad-port NIC is used, the PS-S200 can accommodate up to three LANs.

PS-S400 can have dual- or quad-port NICs installed in Slots 3 and 4. If these slots are populated with quad-port NICs, the PS-S400 can accommodate up to five LANs.

PS-S500 supports installation of NICs in Slots 1-4. If Slots 3 and 4 are each populated with quad-port cards, up to six LANs can be accommodated on a single PS-S500.

Additional Non-Inline Topology Examples

Unlike the previous topology examples, a non-inline deployment using a network tap does not require the PacketShaper to be in watch mode. In this topology, the PacketShaper is connected to the network through a tap. A network tap is a transmit-only hardware device, placed inline, that provides a permanent access port for passive network monitoring. When a PacketShaper is connected to a tap, it receives the same traffic as if it were located directly on the wire.

The PacketShaper's INSIDE and OUTSIDE ports are connected to the tap – and **watch mode is not enabled**. The tap will copy traffic in both inbound and outbound directions.

Since the tap does not pass traffic back to the network, you must use the SYS MGMT port to manage the unit.

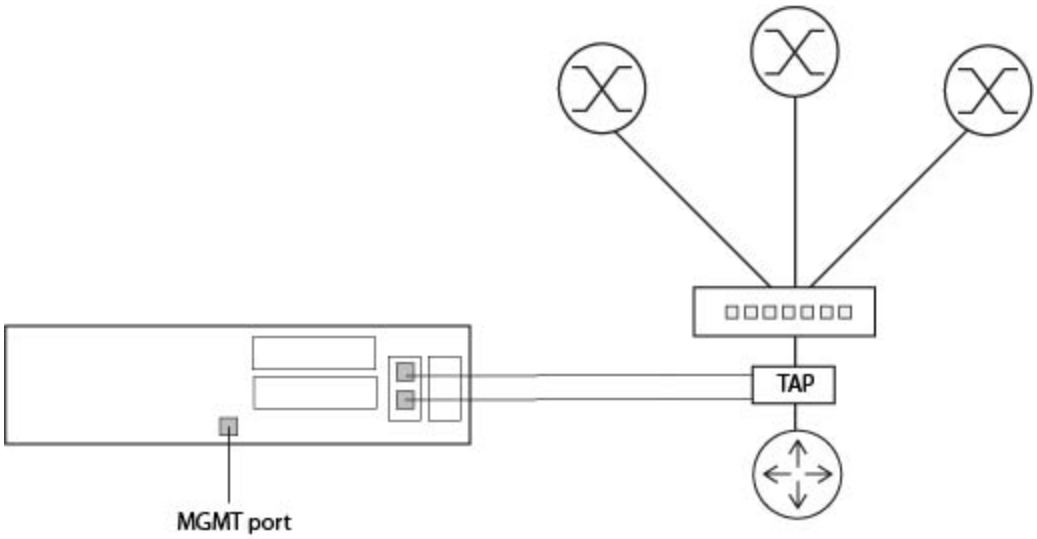

The unit's address, router, and management port connection must be on the same subnet.

Capabilities

A PacketShaper in watch mode includes the following features:

- Is compatible with SPAN ports, mirrored switch ports, and hubs connecting multiple routers
- Can monitor traffic to/from up to 256 WAN routers
- Can monitor traffic from different network segments (if NICs are installed)
- Can be managed from a separate network than the one being monitored
- Can be connected to the SPAN port via the PacketShaper's INSIDE or OUTSIDE port

Limitations

Consistent with its role as a passive monitor, a PacketShaper deployed non-inline with a network tap or in watch mode will not perform traffic shaping; any existing policies and partitions will be ignored.

Configuration Suggestions

For non-inline deployments where a PacketShaper monitors traffic on more than one LAN segment (such as Example 2) you may want to define subnet- or IP address-based traffic classes for each LAN in order to track traffic statistics for each LAN.

Scaling Considerations

The PacketShaper for a non-inline deployment needs to accommodate the link's capacity, anticipated maximum number of concurrent flows, and anticipated number of users. If you are monitoring traffic on more than one LAN, choose a model based on total WAN-link capacity, not the needs of one particular LAN. Be sure to verify if the model that suits your capacity needs supports your desired NIC. Consult the Configuration Limits.

Access the CLI with a Browser

If you are using the PacketShaper browser interface and need to enter a few commands from the command line, you can bring up the Command Interpreter in your browser window.

1. In your browser's **Address** field, type the unit's IP address followed by **/cli.htm** for example, *http://10.10.10.10/cli.htm*.

The Command Interpreter screen appears in your browser window.

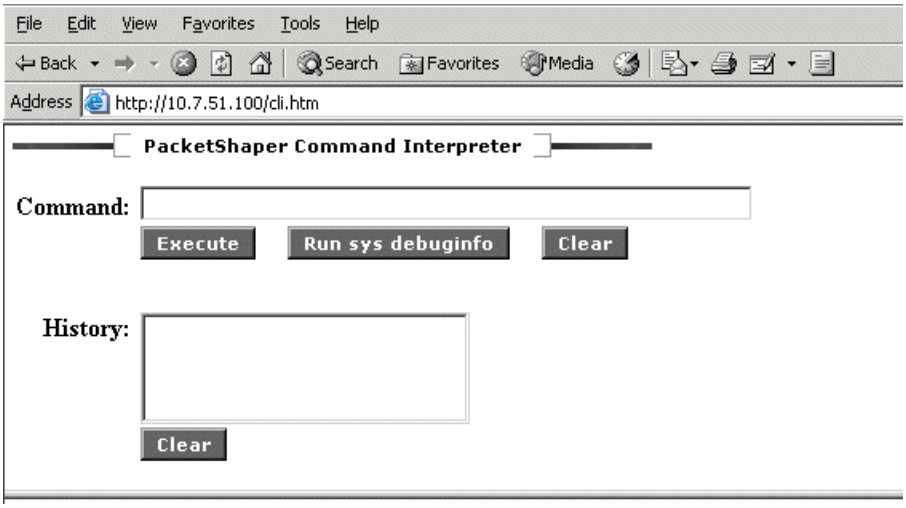

2. Type a command in the **Command** field and click **Execute** (or press Enter).

The command output appears in the bottom window pane, and the command you just entered is listed in the **History** box. To reexecute a command that appears in the **History** box, you can click the command and then click **Execute**.

3. To return to the PacketShaper browser interface, click the browser's **Back** button.

The Command Interpreter has a few limitations you should keep in mind. First, you cannot issue interactive commands that require user input or confirmation. Second, you should not issue a command before the previous command has finished processing.

Access the CLI Using a Console Connection

If remote-access methods do not work due to network or configuration issues, you can access the unit directly with a null-modem cable.

To access the command-line interface via a serial connection:

1. Using the provided null-modem cable, attach a workstation or PC to the unit s port labeled CONSOLE. This cable offers both 9-pin and 25-pin connectors on each end.

2. Start your terminal emulation program (such as HyperTerminal).

3. Verify that you have configured the program with the following values to communicate with the unit s console serial port:

9600 bps, 8 data bits, 1 stop bit, no parity, hardware flow control

If you are using a modem connected to the serial port, the modem must be set to: 9600 bps, 8 data bits, 1 stop bit, no parity, auto-answer (usually ATH1 in the standard Hayes command set), and DTR always on (usually a DIP switch setting). Check the modem manual for details.

4. Power on the unit, if you have not already done so. If the unit was already turned on, you will need to press Enter several times to make the connection.

5. When prompted, enter the user name (**touch** or **look**) and password.

When you successfully connect, you will see the unit's command-line prompt, for example *PacketShaper#*.

Access the CLI Using Remote Login

You are free to choose any remote login utility that is available for your operating system, such as SecureCRT for Windows or OpenSSH for UNIX operating systems. You can have four concurrent SSH sessions.

> **Note:** For a router-installed unit, the unit must be installed on the network between your workstation and the WAN access router. For a server-installed unit, you must log in from the server.

To access the PacketShaper command-line interface with a remote login utility:

1. Log into the unit using its IP address, for example: **ssh 10.10.1.100**. When you connect successfully, you will be prompted to log in.
2. Enter the user name (**touch** or **look**). Or, if RADIUS or TACACS is enabled, type your RADIUS or TACAS user name and press Enter.
3. Type the password and press Enter.

When you successfully connect, you will see the unit's command-line prompt, for example *PacketShaper#*.

Access the CLI with the Quick CLI Commands Utility

The Multi-Class Quick CLI Commands utility allows you to execute CLI commands for multiple classes in one operation. For example, you can apply a 100k partition to six different classes. In addition, the utility provides a command line into which you can enter CLI commands, without having to use a remote login utility.

Types of Commands

The utility can issue the following types of commands on multiple classes:

- Policy (show, apply, dscp, precedence, remove)
- Class (show, discover, delete, copy, move, set, owner, note)
- Misc (packetcapture, hostdb)
- Partition (apply, show, remove)

Selecting Multiple Classes

To select the classes you want to apply the command to, use one of the following methods:

- Click the class name in the Available Classes list, then click the Add button to move it over to the Target Classes list. (Or, you can double-click the class name.)
- Hold down Ctrl as you click each desired name in the Available Classes list, and click the Add button to move all selected classes to the Target Classes list. (You can also use the Shift+Click method to select a group of consecutive classes.)
- Click the Add All button to select all classes in the Available Classes list and add them to the Target Classes list.
- Filter the Available Classes list so that it only displays classes that contain the matched filter string. For example, typing **out** in the **Filter** field will match all classes with *out* in their name (including all Outbound classes). Or if you type **http**, all classes with *http* anywhere in the name will be listed. Clicking the ***Clear*** button will reset the Available Classes list after it has been filtered. Once the list is filtered, you can add classes to the Target Classes list using any of the methods described above.
- If you accidentally add a file to the Target Classes list, you can select it and use the Remove or Remove All buttons.

Running a Quick Command on a Group of Classes

1. Click the **Info** tab.

2. Click **quick commands**. The Multi-Class Quick CLI Commands utility opens. show screen

3. Select the classes in the Available Classes list, and add them to the Target Classes list. (See Selecting Multiple Classes above.)

4. Select a command from the Quick Command drop-down lists. The command appears in the **Command** field (for example, *policy remove <class>*).

> **Note:** The *<class>* that appears in the command is a variable into which each class on the list will be substituted when the command is executed.

5. If you need to modify the command, click the **Enable any CLI Command** checkbox and edit the command as necessary.

6. Click **Run**. The results (output) of the command are displayed in the **Output** window.

7. When you're done, click **Close**.

Running other CLI Commands

You are not restricted to the commands in the Quick Commands drop-down lists – you can enter any non-interactive CLI command, such as **traffic tree**. An interactive command is one that requires user input after the command is entered. If you enter an interactive command, the PacketShaper may have to be power-cycled. Interactive CLI commands known to cause problems include:

- **setup guide**
- **setup reset**
- **class reset**
- **class load**
- **chkdsk**

To run non-interactive CLI commands:

1. Select the **Enable any CLI Command** checkbox.

2. In the **Command** field, type the non-interactive CLI command.

3. Click **Run**.

Access the Command-Line Interface

The command-line interface (CLI) provides a UNIX-like interface for accessing the PacketShaper software. All of the functions available via the browser interface are also accessible with commands. In addition, a number of CLI commands support diagnostic tasks that are not incorporated in the browser interface.

There are three ways to access the CLI:

- Remote login utility
- Console connection
- Browser interface

Agent Host Lists

The majority of agents in the host category have two host lists associated with them: exception and violating. In the exception host list, you can add the IP addresses of the hosts you don't want to trigger the agent (such as servers). The violating host list contains the IP addresses of the hosts that exceed the threshold set for the agent; adaptive response automatically adds hosts to this list when they are in violation. For example, for the Failed Flow Ratio agent, hosts that exceed the RatioThreshold will automatically be added to the violating host list.

The agent templates that use host lists are:

- NFPM Side Unknown
- NFPM Failed Flow
- Failed Flow Ratio
- Host Info Variables

Host Exception Lists

If you want a single exception host list to apply to all host agents, you can add exception hosts to the default host list that is automatically created for you: exceptionHosts. If you want different host lists for different agents, you will need to create host lists with unique names.

1. Access the command-line interface.

2. Type hl show to see the host list names.

3. To create a new host list, use the hl new command. For example:

```
hl new NFPMexception
```

4. To add hosts to an existing host list, use the hl add command. You can add multiple addresses or subnets by separating each with a space. For example:

```
hl add exceptionHosts 192.21.18.162 192.21.18.165 192.21.18.169
```

5. When creating or editing the agent, make sure to specify this host list name for the ExceptionHosts parameter.

Violating Host Lists

Adaptive response automatically creates a violating host list using the name supplied in the agent's ViolatingHosts parameter field. This list is named violatingHosts unless you change the name. Any hosts that exceed the agent's thresholds will be added to this host list. The maximum number of hosts in the violating hosts list is 1000; after the limit is reached, no additional hosts will be added to the list. Adaptive response automatically removes a host from the violator host list if it doesn't have activity for five minutes.

> **Note:** All host agents have the same default violating host list name (violatingHosts). If you want to maintain unique host lists for each agent, make sure to change the name when defining the agent.

Restricting Bandwidth for Violating Hosts

In order to apply restrictive bandwidth controls on the violating hosts, you may want to create traffic classes that are based on the violating host list names.

For example, to create Inbound and Outbound classes in the command-line interface, use commands similar to the following:

```
class new inbound hostviolators outside list:violatinghosts inside list:violatinghosts
 class new outbound hostviolators outside list:violatinghosts inside list:violatinghosts
```

In the browser interface, create an Inbound class and select the violating host list name from the Host List drop-down lists (Inside and Outside). Repeat for Outbound.

After you have created these classes, you can apply appropriate policies and/or partitions to restrict the violating hosts' bandwidth.

Application-Specific Matching Rule Criteria

Overview

Some services can be further classified with specific criteria, enabling deeper classification within the service type. You can specify the application criteria when creating classes in Advanced UI or in the Sky UI.

The applications listed below can be further classified by attributes that are specific to the application.

- Citrix-ICA
- DCOM
- DICOM
- FTP-Data-Clear and FTP-Data-Clear-IPv6
- HTTP, HTTP-IPv6, and SOAP-HTTP
- HTTP-Tunnel and HTTP-Tunnel-IPv6
- ICMP
- NNTP-Clear
- Oracle-netv2
- PostgreSQL
- ProxySG-ADN and ProxySG-ADN-IPv6
- RTCP-I
- RTP-I
- SMTP-Clear and SMTP-Clear-IPv6
- SSL and SSL-IPv6
- WAP

If N/A appears in the matching rule's **Criterion** field, application-specific criteria are not available for the selected service. For special considerations, see Specifying Multiple Criteria.

Citrix-ICA Criteria

Citrix-ICA (Independent Computer Architecture) can be further classified in three ways: by published application, by client name, or by priority.

Classifying by Published Application

Citrix-ICA can be subclassified by published application. The application classes can be created either automatically or manually.

Automatic Creation of Published Application Classes

When traffic discovery is enabled, PacketShaper will automatically discover Citrix published applications and create classes for them. For instance, when PacketShaper detects Citrix traffic, a Citrix class will be created, and class discovery will be enabled on that class. Because class discovery is on, PacketShaper will also create a Citrix/Default class and Citrix child classes for any published applications.

If your traffic tree already has a Citrix class, you can manually enable class discovery, whereupon the application-named classes will be automatically created as child classes of Citrix.

Manual Creation of Published Application Classes

To manually create a class for a Citrix-ICA published application, create a class for the Citrix-ICA service, and select **published application** from the **Criterion** drop-down list. You then must supply a specific name (up to 99 characters). The published application name you enter must exactly match the name configured on the Citrix server and client. **show screen...**

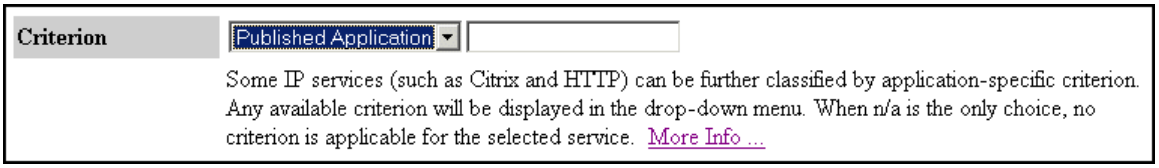

For example, you could create a class for Citrix-ICA traffic that carries the published application PeopleSoft. To create a PeopleSoft class as a child class of the Citrix class, you would specify the following attributes:

- Name: PeopleSoft
- Protocol: IP
- Service: Citrix-ICA
- Location: any
- Criterion: Published Application
- Criterion (second field): PeopleSoft

> **Note:** The entry in this second Criterion field must match the name for the published application defined in the Citrix Program Neighborhood.

Classifying by Client Name

To classify by a Citrix-ICA client name, specify Citrix-ICA as the service, and select **client name** from the **Criterion** drop-down list. You then must supply a specific name (up to 99 characters). The client name you enter must exactly match the name configured on the Windows client. **show screen...**

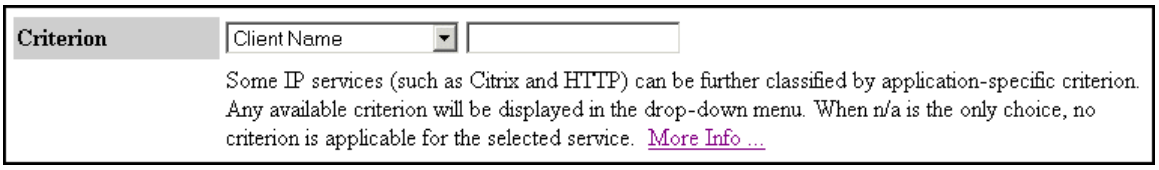

Classifying by Priority

Citrix-ICA has the ability to assign different priority levels to virtual channels within a single ICA flow, and PacketShaper can classify the traffic according to these priority levels. With Citrix-ICA specified as the service, select **priority** from the **Criterion** drop-down list and supply a value from 0 to 3 (with 0 being the highest priority). You may want to create Citrix child classes for each priority level. **show screen...**

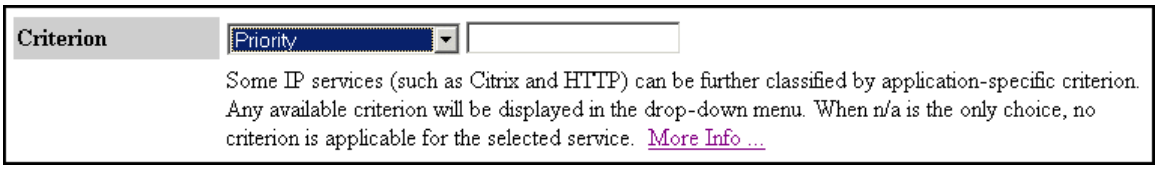

> **Note:** The Citrix priority numbers are used for classification purposes only. Once the traffic is categorized by Citrix priority, you can use PacketShaper policies to manage bandwidth on the class.

DCOM Criteria

DCOM (Distributed Component Object Model) allows client program objects to request services from server program objects on other computers in a network. Using the UUID attribute, you can classify DCOM traffic with a specific Unique Universal Identifier. To do this, create a class based on the DCOM service, and for the **RPC UUID** criterion, enter the UUID. (You can use the class criteria commands in the CLI to determine what value to enter for the UUID.)

DICOM Criteria

DICOM (Digital Imaging and Communications in Medicine) is the global industry standard for transfer of radiological images such as MRI, CT, PET, ultrasound, and mammography. You can create classes for specific DICOM client or server applications using the **Server Title** or **Client Title** criteria. With DICOM selected as the service, select either **Server Title** or **Client Title** from the **Criterion** drop-down list field and enter the name of the server or client. **show screen...**

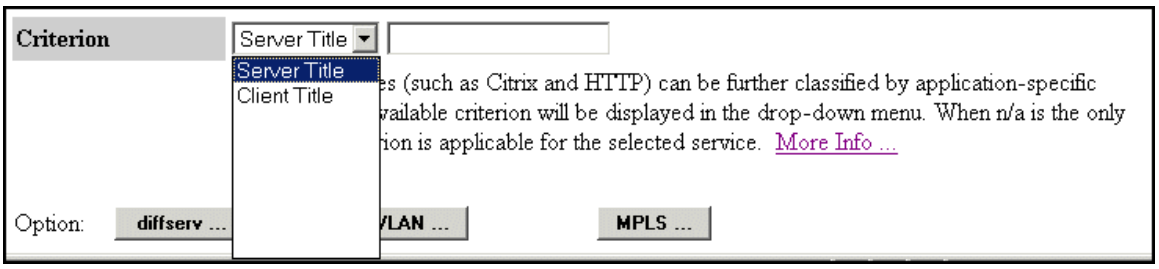

> **Tip:** You can use the class criteria commands in the CLI to determine what value to enter for the criterion.

> **Note:** PacketShaper can automatically discover DICOM client or server applications when DICOM has traffic discovery enabled within the class.

FTP-Data-Clear and FTP-Data-Clear-IPv6 Criteria

To classify FTP downloads by file extension or filename, select FTP-Data-Clear or FTP-Data-Clear-IPv6 as the service and use the **File Name** criterion. You can specify an exact filename or use wildcards (* and ?). For example, to classify FTP downloads of MP3 files, you can specify **.mp3* as the **File Name** criterion. **show screen...**

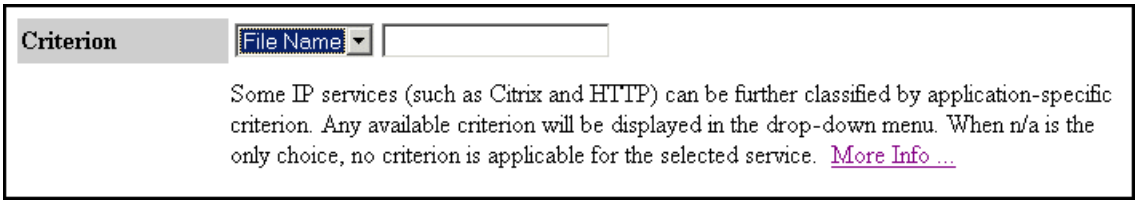

HTTP, HTTP-IPv6, and SOAP-HTTP Criteria

When HTTP, HTTP-IPv6, or SOAP-HTTP is listed as the service in a matching rule, the following attributes can be specified in the **Criterion** field to qualify the HTTP service type:

- Host DNS Name or IP Address
- URL
- Content Type (for example, image/gif)
- Web Browser or User Agent

show screen...

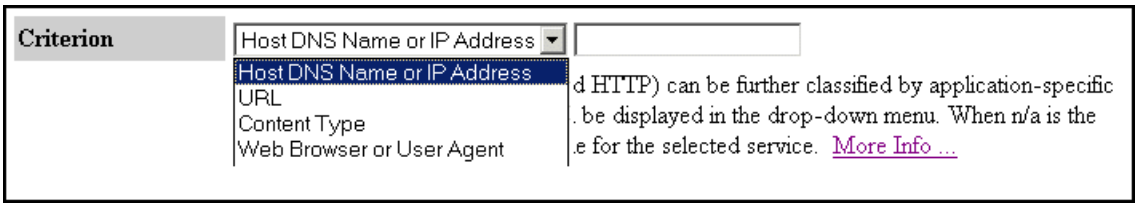

Host DNS Name or IP Address Classification

If you select **Host DNS Name or IP Address** as your web criterion, you can specify the DNS name or IPv4 address of a website in order to control access to the site. The wildcard characters (***** for a group of characters and **?** for single characters) are supported.

> **Note:** You cannot enter an IPv6 address as a web criterion.

URL Classification

To classify traffic for a specific web page, select **URL** from the **Criterion** drop-down list in the matching rule and specify a URL according to the syntax rules shown below.

- Specify only the portion of the URL that follows the server hostname (http://hostname). The hostname goes into the **Name** field (in the Inside or Outside column in the window, depending on where the web server resides).

  For example, to specify a remote website *http://www.mycompany.com/FavoriteStuff*, enter **www.mycompany.com** into the **Name** field of the Outside column and enter **/FavoriteStuff** in the URL **Criterion** field.

- The wildcard characters (***** for a group of characters and **?** for single characters) are supported.

  For example, you can configure a traffic class to look for a certain type of graphics file by specifying **.gif* as the URL.

- Slashes are *not* treated as special characters.

  To match a specific URL, use the full pathname, *including* the leading slash.

- Only one URL can be specified per matching rule.

> **Note:** PacketShaper examines only the first 128 bytes of the URL in a flow.

Content Type Classification

Content types describe specific types of web objects, such as JPEG images or HTML text. To identify content types, use the command-line interface:

1. Access the CLI.

2. Use the **class criteriatrack** command to identify the type of web objects traversing the link. For example:

class criteria track /inbound/http web content-type

3. Generate some web traffic.

4. Use the **class criteria recent** command to show recent values for a class. For example:

class criteria recent /inbound/http

```
Traffic Class: /Inbound/HTTP
Application: Web
Attribute: content-type (Content type)
Recent Attribute Values (most recent first)
---------------------------------------------------------------
1. text/html
2. image/gif
3. text/plain
4. image/jpeg
```

5. Turn off tracking when you are done collecting data. For example:

class criteria track /inbound/http off

When specifying the Content Type criterion, enter the content type as listed in the **class criteria recent** command output (see step 4 above).

When classifying HTTP traffic by content type, if the content type's value is also a service, the traffic will be classified as the service, not as HTTP. For example, the HTTP content type *video/mpeg* is automatically classified as the service MPEG-Video, even if there is a manually created HTTP class with the criterion *web:content-type:video/mpeg*.

Web Browser (User Agent) Classification

The user agent is a string that identifies what web client software is being used. The content of this field is at the discretion of the browser developer.

To identify user agents, use the command-line interface:

1. Access the CLI.

2. Use the **class criteria track** command to identify the web clients being used on the link. For example:

class criteria track /inbound/http web user-agent

3. Generate some web traffic.

4. Use the **class criteria recent** command to show the recent user-agent values for a class. For example:

class criteria recent /inbound/http

```
Traffic Class: /Inbound/HTTP
Application: Web
Attribute: user-agent (Web browser or user agent)

Recent Attribute Values (most recent first)
---------------------------------------------------------------------
1. Mozilla/5.0 (Windows; U; Windows NT 5.0; en-US; rv:1.4) Gecko/20030624 Netscape/7.1 (ax)
2. Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0)
3. Mozilla/5.0 (Windows; U; Windows NT 5.0; en-US; rv:1.7.5) Gecko/20041107 Firefox/1.0
```

5. Turn off tracking when you are done collecting data. For example:

class criteria track /inbound/http off

In this example, all browsers use *Mozilla* as an attribute value. To differentiate between these browsers, you would need to use something more specific than *Mozilla* as the User Agent criterion. For Microsoft Internet Explorer, you could use *MSIE* as the criterion and for Netscape you could use *Netscape/7.1.*

Additional Information

- Web user agent strings often contain spaces. When specifying a string as a user agent criterion in the Legacy user interface, *donot* enclose the string in quotes–they cause screen display and classification problems. In the CLI, however, you *must* enclose the string in double quotes when the string contains spaces.
- The * wildcard is supported when specifying user agent strings.

Understanding Web Criterion Classification

The nature of web traffic dictates how specific web criteria are classified, as noted in the following list:

- **web:host** classifies on the request direction, and in the response direction only after the flow is reclassified, because there is no host header in the response.
- **web:content-type** classifies on the response direction.
- **Web:user-agent** classifies on the request direction, and in the response direction only after the flow is reclassified.
- **Web:url** classifies on either direction.

HTTP-Tunnel and HTTP-Tunnel-IPv6 Criteria

Traffic that is sent through an HTTP tunnel via an HTTP proxy server on the Internet is classified as HTTP-Tunnel. HTTP-Tunnel traffic can be further classified by host or by port number.

To classify by a particular HTTP proxy server, create a class for the HTTP-Tunnel or HTTP-Tunnel-IPv6 service, select **Host DNS Name or IP Address** for the **Criterion**, and enter the host's DNS name or IPv4 address in the text field. Or, to classify by a specific port number that is using HTTP-Tunnel, create a class for the HTTP-Tunnel service, select **Port** for the **Criterion**, and enter the port number in the text field.

Note: You cannot enter an IPv6 address as a web criterion.

ICMP Criteria

ICMP traffic can be broken down into ten subtypes, based on the ICMP message type field. These criterion identifiers include:

- address_mask
- dest_unreachable echo
- info
- nachi
- param_prob
- quench
- redirect
- time_exceeded
- timestamp
- tunnel

To classify by ICMP, create a class for the ICMP or ICMPv6 service, make sure **ICMP Type** appears in the **Criterion** field, and type one of the above messages in the empty text field. For example, the UNIX ping command generates ICMP echo request packets and listens for the subsequent ICMP echo response packets. To identify ping traffic passing through the unit, create an ICMP class and specify **echo** for the **Criterion**. Or to create a class to identify and control ICMP tunnel traffic, specify **tunnel** for the criterion. **show screen...**

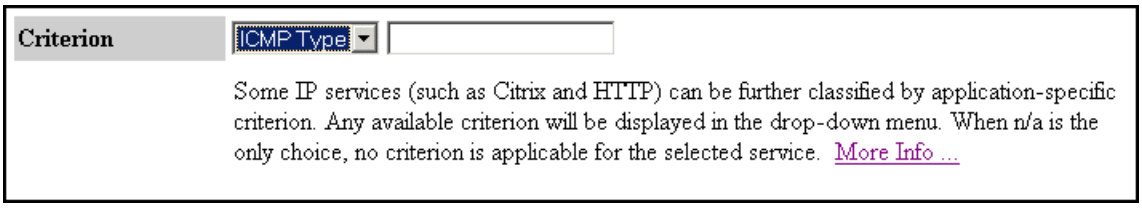

NNTP-Clear Criteria

Network News Transfer Protocol (NNTP) traffic can be further classified by newsgroup name. For example, you can create a class for the Microsoft computer games newsgroup (comp.games.microsoft). To classify by newsgroup name, create a class for the NNTP-Clear service, make sure **Group Name** appears in the **Criterion** field, and type the newsgroup name in the empty text field. You can specify an exact name or use the asterisk wildcard. For example, to classify newsgroup names that have "linux" in the name, you can specify **linux** as the **Group Name** criterion. **show screen...**

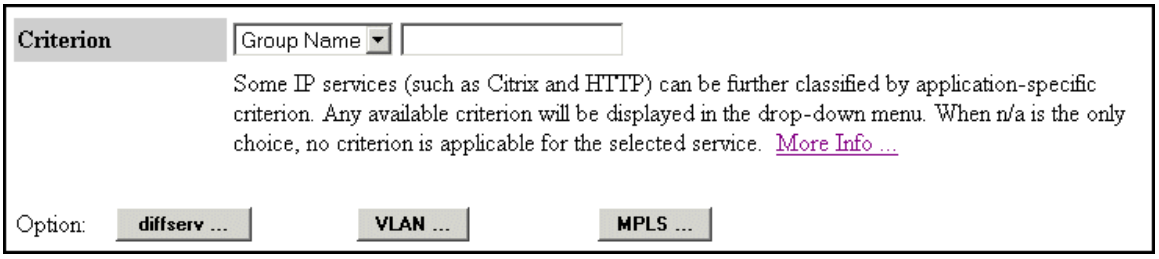

Oracle-netv2 Criteria

Oracle-netv2 traffic can be further classified by database name. PacketShaper looks at the database name that the client requests when it connects to the Oracle listener. This means that you can classify by database name only if you are connecting to the database by name, rather than by requesting the "default database." The database classes can be created either automatically or manually.

Automatic Creation of Database Classes

When traffic discovery is enabled, PacketShaper will automatically discover Oracle database names and create classes for them. For instance, when PacketShaper detects Oracle traffic, an Oracle class will be created, and class discovery will be enabled on that class. Because class discovery is on, PacketShaper will also create an Oracle/Default class and Oracle child classes for any Oracle databases.

If your traffic tree already has an Oracle class, you can manually enable class discovery, whereupon the database-named classes will be automatically created as child classes of Oracle.

Manual Creation of Database Classes

To manually create an Oracle database class, create a class for the Oracle-netv2 service, make sure **Database Name** appears in the **Criterion** field , and type the database name in the empty text field. **show screen...**

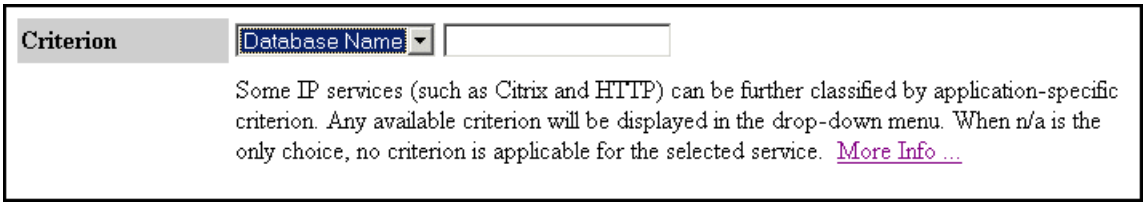

PostgreSQL Criteria

The PostgreSQL freeware SQL database application can be subclassified by database name. The database classes can be created either automatically or manually.

Automatic Creation of Database Classes

When traffic discovery is enabled, PacketShaper will automatically discover PostgreSQL database names and create classes for them. For instance, when PacketShaper detects PostgreSQL traffic, a PostgreSQL class will be created, and class discovery will be enabled on that class. Because class discovery is on, PacketShaper will also create a PostgreSQL/Default class and PostgreSQL child classes for any PostgreSQL databases.

If your traffic tree already has a PostgreSQL class, you can manually enable class discovery, whereupon the database-named classes will be automatically created as child classes of PostgreSQL.

Manual Creation of Database Classes

To manually create a PostgreSQL database class, create a class for the PostgreSQL service, make sure **Database Name** appears in the **Criterion** field, and type the database name in the empty text field. **show screen...**

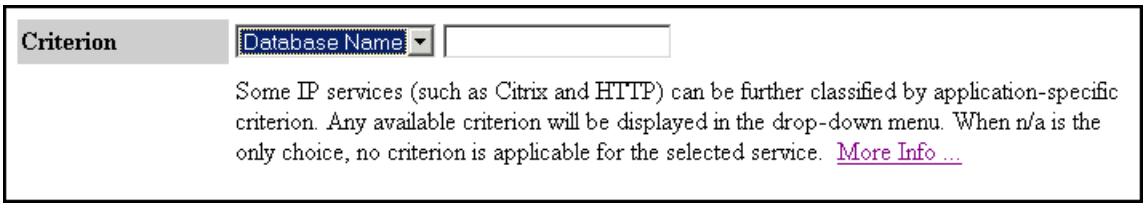

ProxySG-ADN and ProxySG-ADN-IPv6 Criteria

ProxySG ADN tunnels can be subclassified by application or port number (for example, ProxySG-ADN-CIFS or ProxySG-ADN-1666). These application-specific or port-specific classes can be created either automatically or manually.

Automatic Creation of ProxySG-ADN Application/Port Classes

When traffic discovery is enabled, PacketShaper will automatically discover Inbound and Outbound ProxySG-ADN classes and create child classes for each application or port that is used in the ADN tunnel. For instance, when PacketShaper detects ProxySG-ADN traffic, a ProxySG-ADN class will be created, and class discovery will be enabled on that class. Because class discovery is on, PacketShaper will also create a ProxySG-ADN/Default class and child classes for any application or port traffic using an ADN tunnel.

If your traffic tree already has a ProxySG-ADN class, you can manually enable class discovery, whereupon the application- or port-named classes will be automatically created as child classes of ProxySG-ADN.

Manual Creation of ProxySG-ADN Application/Port Classes

To manually create a ProxySG-ADN application or port-based class, create a class for the ProxySG-ADN or ProxySG-ADN-IPv6 service, make sure **App/Port** appears in the **Criterion** field, and type the application name or port number in the empty text field. **show screen...**

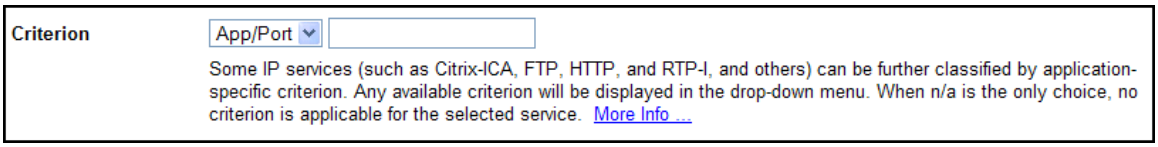

RTCP-I Criteria

Real-time control protocol (interactive) can be further classified by the following criteria. Using these criteria, you can differentiate between different audio and video streams. Or, you could set the encoding name to a value such as *GSM* or *PCMA* so that you can manage streams differently based on how they are encoded.

Attribute	Examples of Values
Encoding Name	G729, GSM, JPEG, PCMA, PCMU **Note:** PCMA and PCMU are both specified in CCITT/ITU-T recommendation G.711. To specify G.711 encoding, enter either **PCMA** (Pulse Code Modulation a-law) or **PCMU** (Pulse Code Modulation mu-law) for the Encoding Name. PacketShaper can automatically discover encoding names when the RTCP class has traffic discovery enabled within the class.
Media Type	**a** (audio), **v** (video)
Clock Rate	Supported values include: 8000, 16000, 44100, 90000

To classify by RTCP, create a class for the RTCP-I service and select one of the options from the **Criterion** drop-down list:

- If you select **Encoding Name** for the criterion, specify the encoding name in the adjacent field. (the Sky UI provides a drop-down list of encoding names; the Advanced UI requires that you enter the value. For a list of encoding names, go to ietf.org and look up RFC 1890.)
- If you select **Media Type**, type **a** for audio or **v** for video in the text field (Advanced UI) or select **audio** or **video** from the drop-down list (the Sky UI).
- If you select **Clock Rate**, type a value such as **8000**, **16000**, **44100**, or **90000** (Advanced UI) or select the value from the drop-down list (the Sky UI).

RTP-I Criteria

This real-time protocol for media streaming can be further classified by the following criteria, based on the RTP standard. In addition to the criteria available for RTCP-I, RTP-I also includes criteria for SIP-based RTP traffic so that you can classify by caller, callee, user-agent, source IP address of the call setup (SIP) flow, or destination IP address of the flow. For example, you can classify all VoIP traffic going through a SIP Gateway by creating an RTP class with criteria that identifies the source or destination IP address of the SIP Gateway.

Attribute	Examples of Values
Encoding Name	**dynamic** (for dynamic codec numbers), **G729**, **GSM**, **JPEG**, **PCMA**, **PCMU** **Note:** PCMA and PCMU are both specified in CCITT/ITU-T recommendation G.711. To specify G.711 encoding, enter either **PCMA** (Pulse Code Modulation a-law) or **PCMU** (Pulse Code Modulation mu-law) for the Encoding Name. PacketShaper can automatically discover encoding-name classes when the RTP class has traffic discovery enabled within the class.
Media Type	**a** (audio), **v** (video)
Clock Rate	Supported values include: 8000, 16000, 44100, 90000
SIP Callee Identifier	+12125551212@server.phone2net.com, 17476004249@10.10.254.56
SIP Caller Identifier	agb@bell-telephone.com, 17476004230@172.21.1.41
SIP User-Agent	*X*Lite*, *Asterisk*PBX*, *3CXVoipPhone** **Note:** The asterisk (*) wildcard is supported. PacketShaper can automatically discover user-agent classes if you enable the enableVoIPUseragentAutoDiscovery system variable.
SIP srcIP Address	207.78.98.18
SIP destIP Address	207.78.98.18

To classify by real-time protocol, create a class for the RTP-I service and select one of the options from the **Criterion** drop-down list:

- If you select **Encoding Name** for the criterion, specify the encoding name in the adjacent field. (the Sky UI provides a drop-down list of encoding names; the Advanced UI requires that you enter the value. For a list of encoding names, go to ietf.org and look up RFC 1890.)
- If you select **Media Type**, type **a** for audio or **v** for video in the text field (Advanced UI) or select **audio** or **video** from the drop-down list (the Sky UI).
- If you select **Clock Rate**, type a value such as **8000**, **16000**, **44100**, or **90000** (Advanced UI) or select the value from the drop-down list (the Sky UI).
- If you select one of the SIP criteria, enter the text string in the text field (Advanced UI and the Sky UI). For example, to match all Motorola models, you can enter **Motorola** for the SIP User-Agent criteria.

Here is an example to identify VoIP user agents:

1. Access the CLI.

2. Use the **class criteria track** command to identify the user agents used in RTP-I. For example:

```
class criteria track /inbound/RTP-I/Default RTP user-agent
```

3. Initiate VoIP calls.

4. Use the **class criteria recent** command to show recent values for a class. For example:

```
class criteria recent inbound/RTP-I/Default

Traffic Class: /Inbound/RTP-I/Default
Application: RTP
Attribute: user-agent (SIP User-Agent)

Recent Attribute Values (most recent first)
 ------------------------------------------------------------------------------
1. *X*Lite*
```

5. Turn off tracking when you are done. For example:

```
class criteria track inbound/RTP-I/Default off
```

Once you have determined the user agent (**X*Lite** in the above example), you can create a class for this type of RTP traffic. Create a class based on the RTP-I service, choose **SIP User-Agent** in the **Criterion** field, and enter the user agent string exactly as it appeared in the output of the **class criteria recent** command.

SMTP-Clear and SMTP-Clear-IPv6 Criteria

Simple Mail Transport Protocol (SMTP) can be sub-classified by the sender's email address in the SMTP message header. You can classify for a specific sender's email address or use wildcards to classify all email senders from a specific domain.

To classify by email sender's name, create a class for the SMTP-Clear or SMTP-Clear-IPv6 service, make sure **Sender Email** is selected for **Criterion**, and enter the sender's email address in the empty text field. The * wildcard may be used to classify email for a specific domain. For example, **@test.com* would match any email sender from the test.com domain.

SSL and SSL-IPv6 Criteria

Some peer-to-peer (P2P) applications allow users to transfer files via SSL on port 443 (https). To limit or block this type of traffic, you can create an SSL class that is based on a particular SSL certificate common name (such as my.loudpc.com or www.redhat.com). To identify common names, use the command-line interface:

1. Access the CLI.

2. Use the **class criteria track** command to identify the certificate common names used in the SSL traffic. For example:

```
class criteria track /inbound/ssl SSL commonName
```

3. Allow a period of time for SSL traffic to be generated.

4. Use the **class criteria recent** command to show recent values for a class. For example:

class criteria recent /inbound/ssl

Traffic Class: /Inbound/SSL
Application: SSL
Attribute: commonName (Common Name)

Recent Attribute Values (most recent first)
---
1. my.loud.pc
2. www.redhat.com
3. optionslink.etrade.com
4. trades1.optionslink.com
5. onlineca.bankofamerica.com
6. onlineid.bankofamerica.com

5. Turn off tracking when you are done. For example:

```
class criteria track /inbound/ssl off
```

Once you have determined the certificate common name, you can create a class for this type of SSL traffic. Create a class based on the SSL or SSL-IPv6 service, choose **Common Name** in the **Criterion** field, and enter the certificate common name exactly as it appeared in the output of the **class criteria recent** command.

WAP

WAP (Wireless Application Protocol) can be sub-classified by server. To do this, create a class based on the WAP service, and for the **URI from WAP Protocol GET Request** criterion, enter the URI of the server. The * wildcard is supported. For example, you can enter **.bluecoat.com* for the criterion.

Specifying Multiple Criteria

A matching rule can contain only one application-specific criterion. To combine multiple criteria, use two traffic classes, a parent and a child. For example, to match HTTP traffic to a URL and a content type of audio/8track, first create a class that matches the URL, then create a child class under this URL class and specify the content type of *audio/8track*.

Bandwidth Allocation

Before assigning rate and priority policies, it's important to understand a few basic rules about how PacketShaper allocates bandwidth.

- Traffic flows with guaranteed rates (that is, rate policies) are allocated bandwidth first. All other traffic competes for the remaining or *excess* bandwidth at its policy priority level.
- Traffic classes that don't have a specific policy assigned are given a priority 3 policy (the default).
- If a rate policy is designated burstable at a certain priority level, any available excess bandwidth is allocated in chunks according to this priority level.
- Priority policies get any excess bandwidth they need/want.

Let's look at some simple examples.

Example 1:

Class1 has a 10k guaranteed Rate policy, burstable at priority 3
Class2 has a Priority 4 policy

Class1, being that it has a Rate policy, will get its 10k of guaranteed bandwidth. Then, Class2, being that it is a Priority policy whose level is higher than the burstable priority level of Class1, gets all the excess bandwidth it needs. If Class2 needs all the excess bandwidth, it gets all the excess bandwidth (although Class1 still gets its guaranteed 10k). If Class2 doesn't use all excess bandwidth, Class1 will get excess bandwidth as needed.

Example 2:

Class1 has a Priority 4 policy
All other classes use the default Priority 3 policy

If Class1 is a bandwidth hog, it will get ALL of the available bandwidth. All other classes will be starved and get nothing because Class1 has a higher priority level than the other classes.

As you can see from the above examples, Priority policies should be used with careful consideration of their implications on other traffic and can potentially cause bandwidth starvation of classes with lower priority levels. In general, Priority policies are appropriate for interactive traffic like TN3270, Telnet , or DNS that is latency-sensitive, doesn't burst, and is small.

BCAAA Troubleshooting

Troubleshooting Tools

If you experience problems with BCAAA and the user awareness feature, you can look at event logs in Windows and on the PacketShaper.

Windows Event Log

To troubleshoot the BCAAA service, look at the information, warning, and error messages in the BCAAA event log on the BCAAA server:

1. On your BCAAA server, right-click on **My Computer** and select **Manage**.
2. Select **Systems Tools > Event Viewer > Windows Logs > Application**.
3. Look for messages with BCAAA listed in the Source column.
4. Double-click the message to see details. For example:

> **Note:** Although some of the BCAAA messages refer to ProxySG, these messages apply to PacketShaper as well.

PacketShaper Event Log

The PacketShaper event log shows detailed warning messages if PacketShaper cannot connect to the BCAAA server or if the BCAAA server times out.

1. Access the PacketShaper with a remote login utility.

2. Type the following CLI command:

   ```
   system event display w
   ```

3. Look for messages that refer to BCAAA. For example, the following messages appear when PacketShaper is unable to connect to the BCAAA server:

   ```
   097 [0364295.1319] W: bannerSetSub: /bcaaaprimary, Primary BCAAA Server, host
   10.9.112.191, connection failed (Cannot connect).
    098 [0364309.9081] W: connect_socket: error connecting socket: Connection refused by
   peer.
    099 [0364309.9081] W: initializeBcaaaServer: bcaaa connection initialization failed
   for primary host 10.9.112.191
   ```

If the BCAAA server times out– perhaps because it's under heavy load or has hung, or if there is a network issue–the log might contain messages such as:

```
TIMEOUT receiving list of users response from BCAAA
```

```
TIMEOUT receiving WSSO auth response from BCAAA
```

When these types of warning messages appear in the event log, the status of the BCAAA server on the User Awareness setup page shows *Failed connection.*

Problem: Can't connect to the BCAAA server

Description of the problem: The status of the BCAAA server on the User Awareness setup page shows *Failed connection.*

Solutions: Check the following:

- Verify that you have specified the correct IP address for your BCAAA server on the PacketShaper User Awareness setup page.

- Verify that the port number you specified on the PacketShaper User Awareness setup page matches the one configured in BCAAA.

- Verify that you created a domain user account for the BCAAA service in the Windows Active Directory.

- Can you ping the BCAAA server from your PacketShaper? If you can't, either the server is offline or you don't have a route from your PacketShaper to your BCAAA server.

- If the Windows server on which BCAAA is running has a firewall, ensure that the BCAAA service (bcaaa.exe) is allowed to listen for new connections on the assigned port (16101 by default).

Problem: Can connect to BCAAA server, but user awareness isn't functioning on PacketShaper

- User search and lookups aren't working
- User names don't appear in Top Talkers/Listeners reports
- User names don't appear on Top Hosts pie chart
- Users are not able to sign on with Windows SSO

Description of the problem: The status of the BCAAA server on the User Awareness setup page shows *In Use*, but the **setup bcaaa server-test** command is not able to find a user associated with the specified IP address; the results message displays *The user name could not be determined for this IP*. In addition, the User search functionality does not display any user names, and the Top Talker/Listeners reports and Top Hosts pie charts don't show any user names.

Solution 1: PacketShaper uses Windows Single Sign On (SSO) with Domain Controller Query (DCQ). If DCQ is not enabled or if some domain controllers are not being queried, users will experience problems signing on and the user awareness feature will not function in PacketShaper.

- The Domain Controller Query option may be disabled in the sso.ini file; this option must be enabled when using BCAAA with PacketShaper. To see if this is the problem, view the Windows Event Log and look for the following BCAAA message: *Domain Controller Querying was requested, but it was not configured in BCAAA*. To enable DCQ, you need to edit the sso.ini file, and remove the semi-colon at the beginning of the DCQEnabled=1 line. The sso.ini file is located in the BCAAA installation directory.

- BCAAA must query all Domain Controllers that will process user logins. Failure to query all necessary Domain Controllers is a common problem. If BCAAA isn't querying the right domain controllers, SSO will not work for some users, or may work intermittently. By default, BCAAA queries all DCs in the forest. If you don’t want to query all domain controllers, you can go to the DCQDomainControllers section of the sso.ini file, and list the domain controllers you do want to query.’ But you need to avoid being too restrictive with this list; filtering out too many DCs is a common mistake.

Solution 2: Another possible reason for users not being able to sign on is if a virus scanner is scanning the IP-to-User table files and locks the files, preventing BCAAA from writing to them. To avoid this problem, exclude the following two files from virus scanning on your BCAAA server:

- dcq_primary_full.sso
- dcq_primary_inc.sso

Browser Settings

PacketShaper 11.9 has been tested with the English versions of the following supported browsers:

- Microsoft Internet Explorer 11
- Mozilla Firefox
- Google Chrome
- Apple Safari

Other browsers and versions may be compatible, but have not been tested.

Make sure that you have installed the latest patches from Microsoft.

Be sure your web browser is configured as follows:

- When using Internet Explorer, you may need to turn on Compatibility View if any of the UI screens don't render properly.
- Set to *accept cookies*, the default setting for most web browsers.
- JavaScript enabled.
- Set Internet Explorer to check for newer versions of stored pages *every visit to the page*. (As you modify PacketShaper configurations in the browser interface, the browser must update the page each time rather than displaying the data already stored in the cache.)
- For optimal formatting of features such as graphs, reports, and data charts, set the screen resolution to 1024x768.
- Pop-up blocking disabled.
- If you are using third-party ad-blocking software, you should configure the software to allow pop-ups from the hostname or IP address of your PacketShaper. If that level of configuration is not possible, you should disable pop-up blocking while managing your PacketShaper.

Examples of Bandwidth-Utilization Analysis

Several applications' bandwidth utilization graphs, accompanied by interpretations, are below:

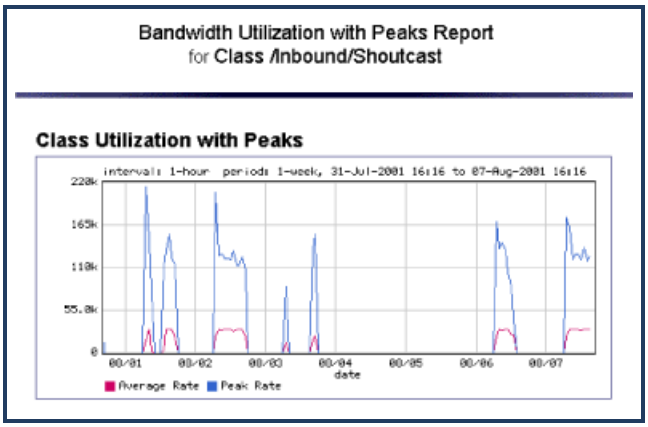

The graph above shows Internet Radio traffic. Notice the relatively constant demand for network resources over time, as indicated by the horizontal pattern of the red, average-rate line. The red line also indicates that average utilization is about 20 Kbps. This could be either a single user listening to a 20 Kbps stream, or multiple users listening to slower streams that total 20 Kbps.

The graph's peaks indicate some burstiness in the traffic that could periodically impact other applications' performance. However, given that this is a one week graph, bursting even once in an hour could cause the peak rate graph to go up. These peaks may just reflect changing channels from one radio station to another.

The graph's pattern shows that sometimes Internet Radio is played for short periods of time; other times it may be on most of the day. For more detailed information you could generate utilization graphs for a day or specific hours of a day.

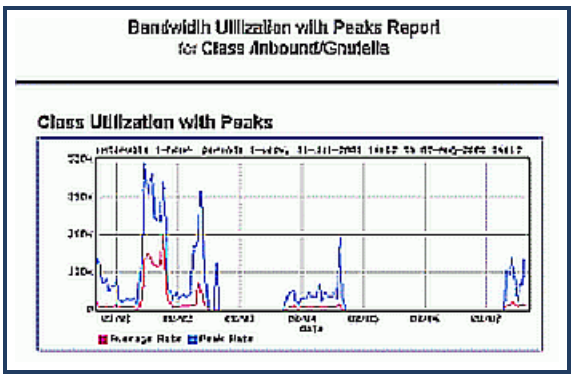

The graph of a music-sharing application provides insights into its use on the network. Note that it's not used on the weekend (8/11 and 8/12). Music-sharing traffic can be significant peaks of 500 Kbps. The number of music-sharing users is probably small. The spikes on the graph probably indicate periods when individual large files are downloaded or when a server on the local network is accessed by a user who downloads multiple files. If there were many music servers there would likely be a more constant stream of traffic. The typical rate usually stays below 100 Kbps and likely has only occasional impact on overall application performance.

The **traffic flow** command could be used to investigate how many concurrent music-sharing sessions there are at the current time. Top Talkers and Top Listeners could identify the top contributors and consumers of this traffic.

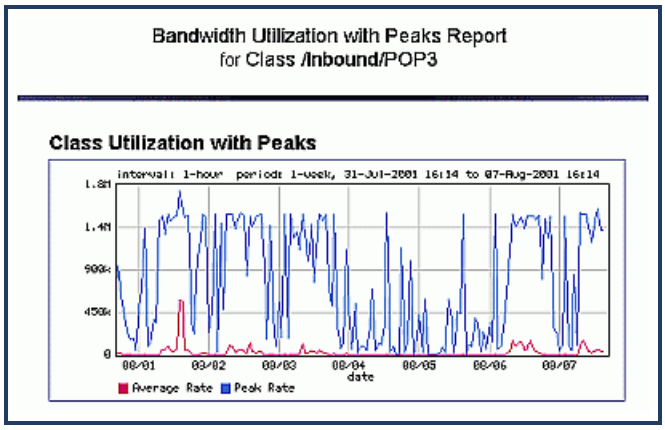

This graph of inbound POP traffic shows downloads from mail servers to individual email clients such as Outlook or Eudora. Like the other applications, this application shows a weekday versus weekend pattern.

Average utilization (red line) is low while the peaks (blue line) remain quite high, meaning mail downloads have bursty and bandwidth-greedy behavior, but don't last long. Since mail software automatically checks for mail on a periodic basis, this traffic is predictably spread out relatively evenly over the course of a day.

During weekdays there are more users, which is reflected in the regular peaks to 1.5 Mbps during business hours. The 1.5 Mbps maximum peaks also suggest that mail server traffic goes through a single T1/E1 line.

As expected, the weekend has either fewer users or fewer PC's turned on and checking for new email. A utilization graph for a single day would probably show business versus non-business hours and patterns that match email use from home in the evening.

The traffic flow command could be used to investigate how many concurrent music-sharing sessions there are at the current time. Top Talkers and Top Listeners could identify the top contributors and consumers of this traffic.

Classification Hints and Examples

If you need to refine the application classes in your traffic tree after PacketShaper's auto-discovery feature created standard classes, the following table of ideas might help. Creative use of matching rule criteria can go a long way in zeroing in on precisely the traffic you have in mind.

Your Classification Needs	How to Accomplish That
I use PVC, VLAN, Citrix, and/or subnet classes. I need to see information on the individual applications under these umbrellas, not just a composite result for all of the applications.	Enable class discovery for these classes. (Note that it is not available for all classes.) Class Discovery creates child classes that each contain a subset of the parent class' traffic.
I need to know which departments use an application.	Using the application's traffic class as a parent, create a child class for each department. The department's IP addresses, subnet, or host list can serve as criteria to identify the department's traffic in the classes' matching rules.
I need to know the top users for an application.	Turn on Top Talkers and Listeners for this class.
I want to group all traffic associated with a single application together (such as MS Exchange or a custom/proprietary application)	If your application is one of those that PacketShaper auto-discovers, then you don't have to create classes yourself. Otherwise, Create a traffic class with a matching rule specifying the IP address of the application's server. Alternatively, you could use a DNS name in the matching rule. Then turn on class discovery for that one class. For example, instead of having a traffic tree that looks like: <...> - Outbound - MsExchange - SAP <...> You could make: <...> - Outbound - MsExchange *(matching rule for DNS name)* - HTTP *(from discovery within a class)* - text *(matching rule for URL=*.htm or mime type)* - graphics *(matchine rule URL=*.jpg or.JPEG etc* - POP3 *(from discovery within a class)* - standard - encrypted - SMTP - Other *(the rest / default)* -SAP <...>

Your Classification Needs	How to Accomplish That
I need more detail on my HTTP web traffic.	Create separate classes for subsets of your HTTP traffic using the variety of criteria in the classes' matching rules. Some ideas include: ■ A unique class for Intranet (versus Internet) traffic using a URL, server(s), and/or travel direction. ■ A unique class for your own customers' web traffic (versus your employees' web browsing) using URL, server(s), and/or direction. ■ Unique classes based on HTTP's mime type to separate HTTP XML, MPEGs, images, text, and so on. That way, for example, you could separate critical XML e-commerce traffic from recreational MPEGs. ■ Unique classes for specific websites using URLs. You could give certain sites or pages preferential treatment (your own sales page, for example). Or you could isolate sites to give poor treatment (sites with questionable content, for example). **Note:** You can use the class criteria track command to identify the type of web traffic on your network.
Which databases is Oracle using?	Create a separate child class under the Oracle class for each database. This would, for example, allow you to track Oracle traffic to the **sales** database separately than that to the **HumanResources** database.

> **Tip:** Dividing classes into distinct child classes also gives you the opportunity to:apply separate policies or partitions and examine performance separately.

PacketGuide solutions that describe creation of a specific type of traffic tree:

- Create an **application-based traffic tree**
- Create a **simple location-based traffic tee**
- Create a **location-based traffic tree with per-location applications**
- Create a **location-based tree with global application**

Command Line Overview

The command-line interface (CLI) provides a UNIX-like interface for accessing the PacketShaper software. All of the functions available via the browser interface are also accessible with CLI commands. In addition, a number of CLI commands support special features and diagnostic tasks that are not incorporated in the browser interface.

Command Usage Conventions

A few basic conventions apply to commands:

- Commands are *not* case sensitive that is, you can use either uppercase or lowercase characters.
- A command can be abbreviated by entering the minimum number of characters required to uniquely distinguish it from other commands. For example, you can type **cl sh** instead of **class show**.
- Command syntax can be verified by typing one of the following:

  help ***<command>***
  or
  <command> **?**

  where *<command>* is the name of the command for which you want help.

- To issue multiple commands from a single command line, separate the commands with a semicolon (;) for example, **`setup show;traffic bandwidth`**. The semicolon is the equivalent of pressing the Enter key.

  > **Note:** When combining multiple commands on one line, do not attempt to run a command file in series with other commands. The **run** command executes a separate task and the other commands in the line may not run in sequence.

- To repeat the last CLI command you entered, type **!!**. To repeat a previous command, type **!*<n>***, where *<n>* corresponds to the sequence of the command in the current remote login or console session. For example, *!5* repeats the fifth command you entered in the current session. Use the **history** command to determine the line number of previous commands. Alternatively, you can scroll through the command history by pressing the up and down arrows. You can also edit previously entered commands.
- *<tclass>* refers to a traffic class name. Include the class' full pathname if it is needed to uniquely identify the class.

  For example, if HTTP appears in both the Inbound and Outbound subtrees, the explicit path is required to identify a specific HTTP class for example, **/inbound/http**.

Typographical Conventions

The following typographical conventions are used for command syntax:

Convention	Description	Example
Boldface	Commands	**class delete web_in**
[Square brackets]	Optional arguments in a command line	**class show [*<tclass>*]**
<angle brackets>	Required arguments for which you will supply a value	**measure dump *<arg>***
Pipe character (|)	The "or" symbol in a command line choose one of the options separated by the | symbol	**setup shaping on|off|bypass**

Editing Previously Entered Commands

If you make a typing mistake in your command, you don't need to retype it you can redisplay the command and edit it. This capability is available via a remote login utility, but not via a direct console connection.

Function	Technique
Display a previously entered command	Press up arrow until the command you want is displayed
Scroll down through the command history	Press down arrow
Move cursor to the left	Press left arrow
Move cursor to the beginning of the line	Press Ctrl+A
Move cursor to the right	Press right arrow
Insert characters	Position cursor and start typing
Delete character to the left of cursor	Press Backspace
Delete all characters on the line	Press Ctrl+U

> **Note:** If the arrow keys aren't working, make sure your remote login utility is emulating VT100 arrows. You may need to enable this option in your client.

Command Line Overview

The command-line interface (CLI) provides a UNIX-like interface for accessing PacketShaper.

Look vs. Touch Access

Users who log in with the touch user name have access to all PacketShaper commands and will see a # sign in the command prompt:

```
PacketShaper#
```

Users who log in with look access will see a > sign in the prompt:

```
PacketShaper>
```

Look users cannot enter commands that modify the PacketShaper configuration.

Command Usage Conventions

A few basic conventions apply to commands:

- Commands are not case sensitive – that is, you can use either uppercase or lowercase characters.
- A command can be abbreviated by entering the minimum number of characters required to uniquely distinguish it from other commands. For example, you can type *set pass* instead of *setup password*.
- Command syntax can be verified by typing one of the following:
  help *<command>*
  or
  ***<command>*?**
  where *<command>* is the name of the command for which you want help.
- To issue multiple commands from a single command line, separate the commands with a semicolon (;) – for example, *setup show;version verbose*. The semicolon is the equivalent of pressing the Enter key.
- You can scroll through the command history by pressing the up and down arrows. You can also edit previously entered commands.

Typographical Conventions

Convention	Description	Example
Boldface	Commands	**class show**
[Square brackets]	Optional argument in a command line	**version [verbose]**
<angle brackets>	Required argument for which you will supply the value	**set images remove *<release_ID>***
Pipe character (|)	The "or" symbol in a command line– choose one of the options separated by the | symbol	**setup discover on|off**

Editing Previously Entered Commands

If you make a typing mistake in your command, you don't need to retype it – you can redisplay the command and edit it. This capability is available via SSH clients and with a direct console connection.

Function	Technique
Display a previously entered command	Press up arrow until the command you want is displayed
Scroll down through the command history	Press down arrow
Move cursor to the left	Press left arrow
Move cursor to the right	Press right arrow
Insert characters	Position cursor with arrows, start typing
Delete character to left of cursor	Press Backspace

> **Note:** If the arrow keys aren't working, make sure your SSH client is emulating VT100 arrows. You may need to enable this option in your client.

Create a Host List

A *host list* contains a set of IP addresses, DNS names, and/or subnets that traffic class matching rules can reference. Host lists let you specify many IP addresses and/or names in a single traffic class matching rule. The addresses in a host list do not need to be contiguous.

The Host List Editor provides an interface for creating, updating, and deleting host lists.

To create a host list:

1. While creating a class or adding a matching rule, select **Host List** in either the Inside or Outside column.

2. Click **edit host lists**. The *Host List Editor* window opens. **show screen...**

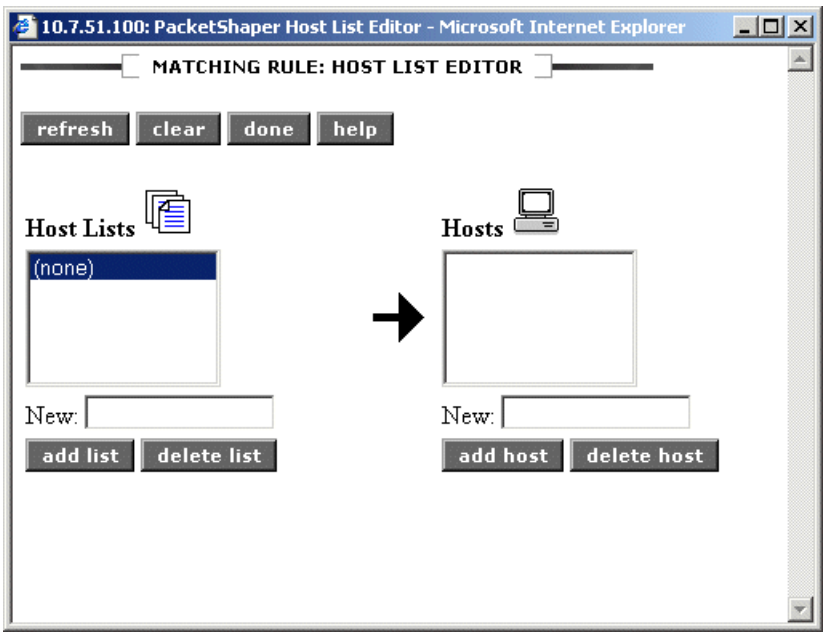

3. Enter a name in the **New** host list field.

4. Click **add list** to generate the new host list.

5. Enter an IP address, subnet, range of IP or subnet addresses, or DNS name in the **New** hosts field. The host can be specified in any of the following ways:

Type of Host	Example
Host IP address	192.168.1.10
Range of IP addresses Use a dash – with no spaces – between the low and high address in the range.	192.168.1.100-192.168.1.200
Address of the subnet; the CIDR number specifies the number of constant bits in the address range	192.168.10.0/24

Type of Host	Example
Range of subnet addresses; the CIDR number specifies the number of constant bits in the address range Use a dash between the low and high address in the range. Spaces are not allowed before or after the dash or slash characters.	192.168.10.0-192.168.20.0/24
DNS name **Note:** Do not use domain names if you will be using the host list with the host sidedness feature.	www.yourcompany.com

> **Note:** The host list editor accepts any addresses and/or names that are syntactically correct. It does not validate the existence of the entries.

6. Click **add host**.

7. Continue adding hosts to complete the list and then click **done** to return to the *Matching Rule* or *Traffic Class* window.

8. To include the selected host list in the matching rule definition, click **apply changes**.

> **Note:** A single Host List Editor session is limited to the creation of 32 host lists. A message window warns when this limit is reached. To initiate a new session to add more host lists, close the *Host List Editor* window and the *Matching Rule* window and then open them both again.

Device Matching Rules in PolicyCenter

When creating a class in PolicyCenter, it is not possible to specify a **Device** matching rule. The **Device** matching rule will be **Device: any** and cannot be changed. This is because NICs may vary between PacketShaper models.

If the class was created in PacketShaper and had a **Device** set to a specific slot, PolicyCenter will recognize the matching rule and will show it as read-only; the **Device** drop-down menu will be greyed out. When a parent configuration has a matching rule based on device, the PacketShaper units subscribed to this configuration will validate the matching rule on devices on the PacketShaper side, instead of the PolicyCenter side. If there is any violation, for example the NIC isn't present, error messages are shown on the PacketShaper side as a configuration error on the class and are displayed in a banner message. Note that the class configuration in PolicyCenter will *not* show the configuration error.

Device

To match traffic that is passing through a specific interface, you can create a class that specifies a physical port pair on the PacketShaper. For example, you can create a class that classifies all outbound traffic through the NIC in Slot 2. You can then turn on class discovery for the NIC class and PacketShaper will auto-discover all the applications, services, and protocols on this NIC. The **Device** field lists only the installed interfaces on your unit. The possible values are:

any (will not classify by device)
slot1
slot2
slot3
slot4

If four-port NICs are installed, you select a port pair (pair1 is on the left, pair2 is on the right):

slot3_pair1
slot3_pair2
slot4_pair1
slot4_pair2

Examples of Efficiency Analysis

Several applications' efficiency graphs, accompanied by interpretations, are below.

The Network Efficiency graph shows the percentage of traffic that was new, initial transmissions (*not* retransmissions). Low efficiency means high retransmission rates, a sign of congested router queues, delays, and time-outs. (**Note:** The formula for efficiency is `bytes - tcp-retx-bytes / bytes`.)

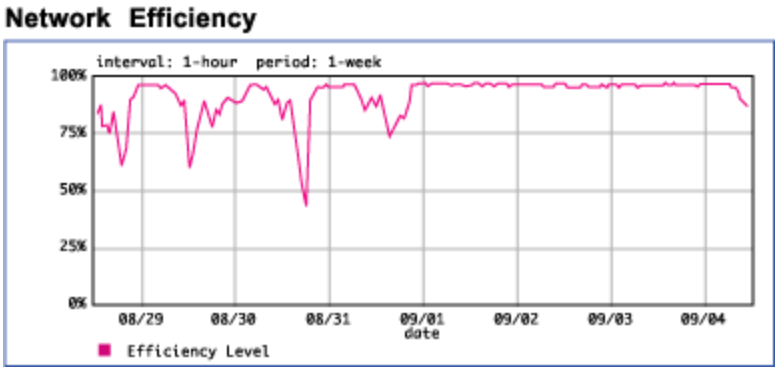

This efficiency graph of a critical ERP application over the course of one week shows what can happen when unmanaged traffic traverses congested lines lots of retransmissions. The drops into the 50 percent range mean that half of active traffic was actually retransmissions. A waste of expensive bandwidth. The second half of the week shows how retransmission can be improved once Shaping is turned on and appropriate policies and partitions are in place. Bandwidth capacity is used more effectively.

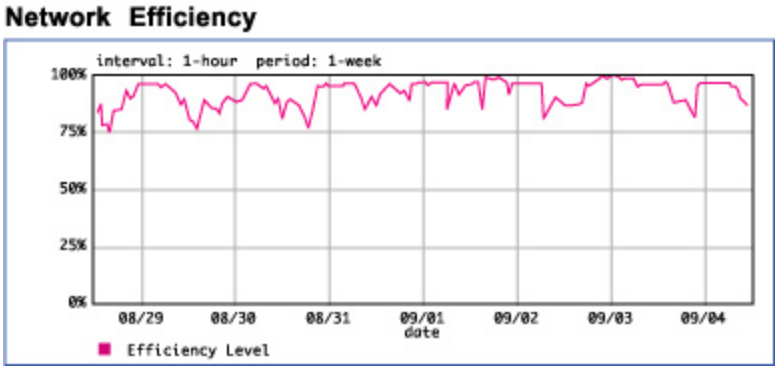

This Inbound POP3 graph's efficiency of about 80 percent shows a 10 to 20 percent retransmission rate for inbound mail traffic. Since email is a mission-critical application, this large percentage of retransmissions should be investigated further to see why so many packets require retransmission. Network efficiency of 95 percent or more is more typical for email on a healthy network with adequate bandwidth. An effective control strategy is appropriate.

FDR Notes

- When using the flow detail records feature, you will need to reset the PacketShaper if you ever change its management IP address. FDRs will not be sent until the unit is reset.

- If your collector is on the LAN and FDR packets are being sent on the INSIDE link, you will not see any measurement data on the NetFlowV5 or FlowRecords classes. Measurement data is recorded on these classes only if FDR packets are sent to the OUTSIDE link. In addition, when FDR packets are sent on the INSIDE link, policies and partitions cannot be activated on these classes.

- FDR records the DSCP value at the time the flow finishes. If you change a DSCP policy in the middle of a flow, the value assigned with this policy is the one recorded in the flow detail record. If you change a DSCP policy for UDP traffic during a time when there are no more packets for the flow (before the flow times out), FDR records the DSCP value that the packet entered the PacketShaper with; in other words, it does not record the value that was assigned to the flow with a DSCP policy.

- Changes made to FDR behavior impact the records emitted to all FDR collectors. For example, you cannot configure your unit to emit intermediate FDRs only to a NetFlow-5 collector, while sending end-of-flow FDRs to a Packeteer-2 collector, such as IntelligenceCenter (which cannot process intermediate FDRs).

FDR Troubleshooting

If your collector isn't receiving flow detail records (FDRs) from your PacketShaper, try the suggestions below to troubleshoot the problem.

Verify Settings

Log in to your PacketShaper, click the **setup** page, and double-check the following:

On the **flow detail record settings** setup page:

- Is the collector's IP address correct?
- Did you specify the correct UDP port?
- Is the collector enabled?

On the **system variables** setup page:

- Does the **Intermediate FDR** setting match the capabilities of your collector?

Log in to your collector and verify the following:

- Is the collector listening?
- Is the collector listening on the same UDP port you have configured in PacketShaper?

Test Network Connectivity

To verify that you can connect to the collector, use the ping command. For example, if the collector's IP address is 10.10.10.10, type **ping 10.10.10.10** at the CLI prompt.

- If you are able to connect to the collector, you will see the message "10.10.10.10 is alive."
- If you are unable to connect, you will see the message "no answer from 10.10.10.10." The server may be down or you may have the wrong IP address for the collector.

You should also confirm that your collector can talk to your PacketShaper. For example, if your PacketShaper address is 192.21.20.190, log into your collector and type **ping 192.21.20.190**.

Verify that FDRs are Being Sent

As PacketShaper sends UDP flow record packets to the collector, it records measurement data for the NetFlowV5 or FlowRecords class. To verify that PacketShaper is emitting FDRs, log into your PacketShaper and go the *Monitor Traffic* screen.

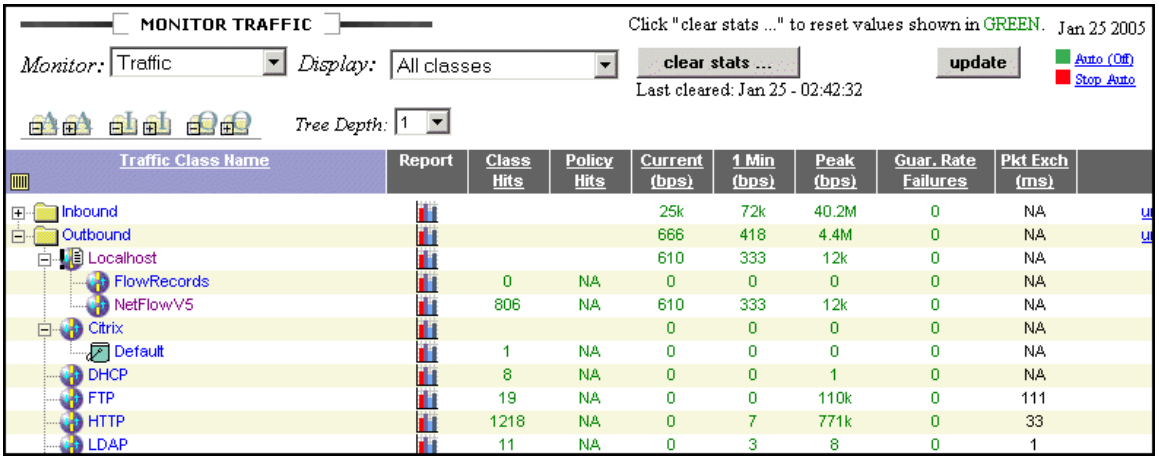

Traffic Class Name	Report	Class Hits	Policy Hits	Current (bps)	1 Min (bps)	Peak (bps)	Guar. Rate Failures	Pkt Exch (ms)
Inbound				25k	72k	40.2M	0	NA
Outbound				666	418	4.4M	0	NA
Localhost				610	333	12k	0	NA
FlowRecords		0	NA	0	0	0	0	NA
NetFlowV5		806	NA	610	333	12k	0	NA
Citrix				0	0	0	0	NA
Default		1	NA	0	0	0	0	NA
DHCP		8	NA	0	0	1	0	NA
FTP		19	NA	0	0	110k	0	111
HTTP		1218	NA	0	7	771k	0	33
LDAP		11	NA	0	3	8	0	1

- Does Outbound/Localhost have a class for NetFlowV5 or FlowRecords? The class is automatically created as soon as a collector is enabled. For example, if a Packeteer-2 collector is defined and enabled, PacketShaper creates a FlowRecords class and this class will track all FDR traffic emitted to the defined Packeteer-2 collector. If a collector is defined, but not enabled, the class will *not* be created. In the example above, Outbound/Localhost/NetFlowV5 and Outbound/Localhost/FlowRecords classes were created.
- Does the *Monitor Traffic* screen show activity in the NetFlowV5 or FlowRecords class? If PacketShaper is sending FDRs and they are being received by a collector, you should see values in the Class Hits, Current (bps), 1 Min (bps), and Peak (bps) columns. In the example above, the NetFlowV5 class does show active data, indicating flow records are being emitted and received; the FlowRecords class does not show any activity (FDRs are not being sent, perhaps because the collector is misconfigured).

Note: If your collector is on the LAN and FDR packets are being sent on the INSIDE link, you will not see any measurement data on the NetFlowV5 or FlowRecords classes. Measurement data is recorded on these classes only if FDR packets are sent to the OUTSIDE link.

Host Agent Templates

The *Hosts* category of agents uses information from PacketShaper's host database. These agents are useful for identifying hosts that are using too much bandwidth or that may be attacking your network, for example spoofing and SYN attacks.

Agent Template Name	Default Agent Name	Description
High Bandwidth Host	na	An agent based on the High Bandwidth Host template monitors hosts, evaluating whether any one is using too much bandwidth, and can help prevent any one host from consuming too much bandwidth. The High Bandwidth Host agent will not return correct values if the unit has more than one instance of this agent type, or if the Quota Bandwidth Host agent is also enabled. The agent tracks hosts *sending* and/or *receiving* an excessive amount of traffic, allowing you to find hosts that are downloading data at high levels. **Agent Parameters** ViolatingHosts: Name of host list. Adaptive response will automatically add violating hosts to this host list. For more information about using violating host lists, see Violating Host Lists. **Note:** All host agents have the same default violating host list name (violatingHosts). If you want to maintain unique host lists for each agent, make sure to change the name when defining the agent. ExceptionHosts: Name of host list that contains hosts to be excluded. You should edit this host list and add the IP addresses or subnets of any hosts (such as servers) that you don't need to monitor. See Host Exception Lists. RedThreshold: An unacceptable percentage of bandwidth consumed by a single host GreenThreshold: An acceptable percentage of bandwidth Default values for the parameters are as follows: ■ RedThreshold >10% of bandwidth on either link ■ GreenThreshold < 5% of bandwidth on either link ■ Yellow 5%-10% of bandwidth on either link ■ Evaluation interval: 1 minute **Action File Variables** $host-ip, $direction, $avg-bps, $violatingHosts, $exceptionHosts **Note:** The $host-ip, $direction, and $avg-bps variables are valid in red action files only. Example

Agent Template Name	Default Agent Name	Description
NFPM Side Unknown	Spoofing - Server, Spoofing - Client	Agents based on the New Flows Per Minute (NFPM) Side Unknown template detect hosts that may be spoofing. Spoofing attacks send packets that appear to be from a trusted source by maliciously setting the source and/or destination to false addresses. The traffic can then gain access to hosts or services that should be secure. A host is considered to be in violation if it exceeds the defined number of new flows per minute. **Agent Parameters** Side: client or server SideThreshold: Number of new flows per minute **New Flows Per Minute** The number of new flows initiated from a host (in the case of a client host) or to a host (in the case of a server) during a one-minute period. ViolatingHosts: Name of host list. Adaptive response will automatically add violating hosts to this host list. For more information about using violating host lists, see Violating Host Lists. **Note:** All host agents have the same default violating host list name (violatingHosts). If you want to maintain unique host lists for each agent, make sure to change the name when defining the agent. ExceptionHosts: Name of host list that contains hosts to be excluded. You should edit this host list and add the IP addresses or subnets of any hosts (such as servers) that you know are not spoofing. See Host Exception Lists. RedThreshold: An unacceptable number of violating hosts; a host is in violation if it exceeds the defined number of new flows per minute (SideThreshold) GreenThreshold: An acceptable number of violating hosts Default values for the parameters are as follows: ■ SideThreshold > 100,000 new flows per minute ■ ViolatingHosts = violatingHosts (name of host list) ■ ExceptionHosts = exceptionHosts (name of host list) ■ RedThreshold >= 1 violating hosts ■ GreenThreshold <= 0 ■ Evaluation interval: 1 minute **Note:** In order for the agents to work properly, you must set the

Agent Template Name	Default Agent Name	Description
		SideThreshold to a value that is appropriate for your network traffic patterns. **Action File Variables** $side, $sideThreshold, $violatingHosts, $exceptionHosts Example

Agent Template Name	Default Agent Name	Description
NFPM Failed Flow	na	An agent based on the New Flows Per Minute (NFPM) Failed Flow template detects hosts that may be SYN attacking. It identifies hosts that have failed flows during the evaluation interval AND that have new flows per minute that exceed the FlowsThreshold. **Agent Parameters** FlowsThreshold: Number of new flows per minute **New Flows Per Minute** The number of new flows initiated from a host (in the case of a client host) or to a host (in the case of a server) during a one-minute period. Side: The traffic direction to be monitored: inside, outside, or both ViolatingHosts: Name of host list. Adaptive response will automatically add violating hosts to this host list. For more information about using violating host lists, see Violating Host Lists. **Note:** All host agents have the same default violating host list name (violatingHosts). If you want to maintain unique host lists for each agent, make sure to change the name when defining the agent. ExceptionHosts: Name of host list that contains hosts to be excluded. You should edit this host list and add the IP addresses or subnets of any hosts (such as servers) that you know are not syn attacking. See Host Exception Lists. RedThreshold: An unacceptable number of violating hosts; a host is in violation if it has any failed flows AND exceeds the defined number of new flows per minute (FlowsThreshold) GreenThreshold: An acceptable number of violating hosts Default values for the parameters are as follows: ■ FlowsThreshold > 100,000 new flows per minute ■ Side = both ■ ViolatingHosts = violatingHosts (name of host list) ■ ExceptionHosts = exceptionHosts (name of host list) ■ RedThreshold >= 1 violating hosts ■ GreenThreshold <= 0 ■ Evaluation interval: 1 minute **Notes:**

Agent Template Name	Default Agent Name	Description

- In order for the agents to work properly, you must set the FlowsThreshold to a value that is appropriate for your network traffic patterns.
- If you set the Side parameter to Inside or Outside, the agent will score only when the PacketShaper sees at least one complete connection for a host.

Action File Variables

$flowsThreshold, $direction, $violatingHosts, $exceptionHosts

Example

Agent Template Name	Default Agent Name	Description
Failed Flow Ratio	na	Agents based on the Failed Flow Ratio template detect hosts that have a high ratio of failed flows compared to new client flows per minute. **Agent Parameters** RatioThreshold: Percentage of new client flows per minute that are failed flows Side: The traffic direction to be monitored: inside, outside, or both ViolatingHosts: Name of host list. Adaptive response will automatically add violating hosts to this host list. For more information about using violating host lists, see Violating Host Lists. **Note:** All host agents have the same default violating host list name (violatingHosts). If you want to maintain unique host lists for each agent, make sure to change the name when defining the agent. ExceptionHosts: Name of host list that contains hosts to be excluded. You should edit this host list and add the IP addresses or subnets of any hosts (such as servers) that you know are not syn attacking. See Host Exception Lists. RedThreshold: An unacceptable number of violating hosts; a host is in violation if it exceeds the RatioThreshold GreenThreshold: An acceptable number of violating hosts Default values for the parameters are as follows: ■ RatioThreshold > 100 percent ■ Side = both ■ ViolatingHosts = violatingHosts (name of host list) ■ ExceptionHosts = exceptionHosts (name of host list) ■ RedThreshold >= 1 violating hosts ■ GreenThreshold <= 0 ■ Evaluation interval: 1 minute **Notes:** ■ In order for the agents to work properly, you must set the RatioThreshold to a value that is appropriate for your network traffic patterns. ■ If you set the Side parameter to Inside or Outside, the agent will score only when the PacketShaper sees at least one complete connection for a host.

Agent Template Name	Default Agent Name	Description
		Action File Variables $ratioThreshold, $direction, $violatingHosts, $exceptionHosts Example

Agent Template Name	Default Agent Name	Description
Host Info Variables	Syn Attack - Failed Flow	With the Host Info Variables template, you can select a variable to monitor hosts: Current Connections, New Flows Client, New Flows Server, Failed Flows. **Agent Parameters** VariableName: Current Connections, New Flows Client, New Flows Server, Failed Flows FlowsThreshold: Number of new flows or connections per minute Side: The traffic direction to be monitored: inside, outside, or both ViolatingHosts: Name of host list. Adaptive response will automatically add violating hosts to this host list. For more information about using violating host lists, see Violating Host Lists. **Note:** All host agents have the same default violating host list name (violatingHosts). If you want to maintain unique host lists for each agent, make sure to change the name when defining the agent. ExceptionHosts: Name of host list that contains hosts to be excluded. You should edit this host list and add the IP addresses or subnets of any hosts (such as servers) that you know are not syn attacking. See Host Exception Lists. RedThreshold: An unacceptable number of violating hosts; a host is in violation if it exceeds the defined number of flows or connections (FlowsThreshold) GreenThreshold: An acceptable number of violating hosts Default values for the parameters are as follows: ■ FlowsThreshold > 100,000 (connections, new flows per minute, or failed flows) ■ Side = both ■ ViolatingHosts = violatingHosts (name of host list) ■ ExceptionHosts = exceptionHosts (name of host list) ■ RedThreshold >= 1 violating hosts ■ GreenThreshold <= 0 ■ Evaluation interval: 1 minute **Notes:** ■ In order for the agents to work properly, you must set the FlowsThreshold to a value that is appropriate for your network

Agent Template Name	Default Agent Name	Description

traffic patterns.

- If you set the Side parameter to Inside or Outside, the agent will score only when the PacketShaper sees at lease one complete connection for a host.

Action File Variables

$variableName, $direction, $flowsThreshold, $violatingHosts, $exceptionHosts

Example

Agent Template Name	Default Agent Name	Description
Quota Bandwidth Host	na	An agent based on the Quota Bandwidth Host template monitors hosts, evaluating whether any one is using too much absolute bandwidth, and can help prevent any one host from consuming too much bandwidth for a specified interval (such as a day). The Quota Bandwidth Host agent will not return correct values if the unit has more than one instance of this agent type, or if the High Bandwidth Host agent is also enabled. The agent tracks hosts *sending* and/or *receiving* an excessive amount of traffic, allowing you to find hosts that are downloading data at high levels. **Agent Parameters** HostUsageThreshold: Maximum number of bytes a host may use during the interval. Threshold can be between 0 and 2,000,000,000 bytes. HostUsageMonitorInterval: Interval to be monitored (in days). Interval can be between 0 and 30 days. At the end of the interval, the measured bytes for each host is zeroed out and any violating hosts are removed from the ViolatingHosts list. Note: The interval begins when the agent is enabled. To control the time of day at which the bytes are reset (such as at midnight), you'll need to enable the agent at the desired time. Side: The traffic direction to be monitored: inside, outside, or both ViolatingHosts: Name of host list. Adaptive response will automatically add violating hosts to this host list. For more information about using violating host lists, see Violating Host Lists. **Note:** All host agents have the same default violating host list name (violatingHosts). If you want to maintain unique host lists for each agent, make sure to change the name when defining the agent. ExceptionHosts: Name of host list that contains hosts to be excluded. You should edit this host list and add the IP addresses or subnets of any hosts (such as servers) that you don't need to monitor. See Host Exception Lists. RedThreshold: An unacceptable number of violating hosts; a host is in violation if it exceeds the defined number of bytes (HostUsageThreshold) during the interval (HostUsageMonitorInterval). GreenThreshold: An acceptable number of violating hosts Default values for the parameters are as follows: ■ HostUsageThreshold >= 10000000 bytes ■ HostUsageMonitorInterval = 1 day ■ Side = both

Agent Template Name	Default Agent Name	Description

- ViolatingHosts = violatingHosts (name of host list)
- ExceptionHosts = exceptionHosts (name of host list)
- RedThreshold >=1 violating hosts
- GreenThreshold <= 0
- Evaluation interval: 5 minutes

Action File Variables

$bytes, $host-ip, $direction, $violatingHosts, $exceptionHosts

Example

Network Health Agent Templates

Network Health agents monitor the health and activity of traffic on the network link. The specific agents in this category and the action file variables for each agent are described in the table below.

Agent Template Name	Default Agent Name	Description
Traffic Performance	na	Agents based on the Traffic Performance template monitor transmission efficiency for one or more traffic classes and can determine when too many retransmissions are occurring, degrading the performance of the traffic class(es). A change in the status of this agent will create a report showing detailed information on traffic performance for the monitored classes over the most current interval. **Note:** Classes with an ignore policy are not monitored. **Agent Parameters** ClassName: Name of traffic class whose efficiency you want to monitor (for example, */inbound/ftp*); a wildcard may be used to monitor many classes (for example, */inbound** to monitor all classes in the inbound direction) Efficiency: Percent efficiency RedThreshold: An unacceptable number of violating classes; a class is in violation if it drops below the defined Efficiency percentage GreenThreshold: An acceptable number of violating classes Default parameters are as follows: ■ ClassName = * (all classes) ■ Efficiency < 80 percent ■ RedThreshold >=1 violating classes (at least one class had an efficiency less than the defined Efficiency percent) ■ GreenThreshold <=0 violating classes ■ Evaluation interval: 5 minutes **Action File Variables** $className, $efficiency, $class-id, $network-efficiency Example

Agent Template Name	Default Agent Name	Description
Partition Utilization	na	Agents based on the Partition Utilization template monitor the throughput (avg-bps) utilization measured against the capacity of a partition, and can alert you when you are running out of bandwidth for the partition. Note that this agent measures the utilization based upon the maximum size of the partition, not the configured minimum. A change in the status of this agent will create a report showing detailed information on the monitored classes over the most current interval. **Agent Parameters** ClassName: Name of traffic class whose partition you want to monitor (for example, */inbound/ftp*); a wildcard may be used to monitor many classes (for example, */outbound** to monitor all classes in the outbound direction) Utilization: Percent utilization of the partition RedThreshold: An unacceptable number of violating classes; a class is in violation if it exceeds the defined Utilization percentage GreenThreshold: An acceptable number of violating classes Default parameters are as follows: ■ ClassName = * (all classes) ■ Utilization > 90 percent ■ RedThreshold >= 1 violating classes (at least one class had utilization greater than the defined Utilization percent) ■ GreenThreshold <= 0 violating classes ■ Evaluation interval: 60 minutes **Action File Variables** $className, $utilization, $class-id, $avg-bps **Note:** The $class-id and $avg-bps variables are valid in red action files only. Example

Unit Health Agent Templates

Unit Health agents monitor the health and activity on an individual PacketShaper. There are four different agent templates in this category; these agents and their action file variables are described in the table below.

> **Note:** You can configure green, yellow, and red thresholds by entering values for just the green and red threshold parameters. The yellow threshold is not a settable parameter—it is the range of values between the green and red parameters, and is shown in the table below just for reference.

Agent Template Name	Default Agent Name	Description
Unit Limits	na	An agent based on the Unit Limits template monitors portions of the system limits of the PacketShaper and will notify you when the capacity of the unit nears its limit. This agent tracks values for the number of classes, matching rules, partitions, and dynamic partitions, and reports a value of 99% when there is only one left of any of these monitored entities. A change in the status of this agent will create a report showing detailed information on unit limits over the most current interval. Each unit or PolicyCenter configuration can have only one instance of a Unit Limits agent. **Agent Parameters** RedThreshold: Percent of system capacity that indicates the unit is approaching its limits; when the RedThreshold is exceeded, the agent will have a red status GreenThreshold: Percent of system capacity that indicates the unit is operating at an acceptable level The default parameter values are: ■ Green < 90% of system capacity ■ Yellow 90-98% of system capacity (implied) ■ Red >= 99% of system capacity ■ Evaluation interval: 15 minutes **Action File Variables** none Example

Agent Template Name	Default Agent Name	Description
System Load	na	An agent based on the System Load template monitors the system processing load on the PacketShaper. This agent can alert you when system load becomes too great. Though a very high percentage of system load is not a problem by itself, it often occurs as a symptom of other problems such as a severe traffic spike or malfunction of the unit. A change in the status of this agent will create a report showing detailed information on system health over the most current interval. Each unit or PolicyCenter configuration can have only one instance of a System Load agent. **Agent Parameters** RedThreshold: Unacceptable percent of system processing capacity. These threshold units are similar to the percentage of CPU utilization. GreenThreshold: Acceptable percent of system processing capacity The default parameter values are: ■ Green < 90% capacity ■ Yellow 90%-95% capacity (implied) ■ Red > 95% capacity ■ Evaluation interval: 1 minute **Action File Variables** none Example

Agent Template Name	Default Agent Name	Description
Memory Allocation	na	An agent based on the Memory Allocation template monitors some of the memory limits of the PacketShaper and can alert you about memory failures in TCP, UCB, ACB, host allocation, and classification allocation. These failures which can occur when a heavy traffic load leads to insufficient memory to manage the new flows. In some cases, memory allocation errors can lead to dropped packets. This agent will automatically return to green if there has been no failure for a 60-second period. A change in the status of this agent will create a report showing detailed information on memory allocation over the most current interval. Each unit or PolicyCenter configuration can have only one instance of a Memory Allocation agent. **Agent Parameters** RedThreshold: Unacceptable number of memory allocation failures GreenThreshold: Acceptable number of failures The default parameter values are: ■ Green = 0 failures ■ Yellow = 1 failure (implied) ■ Red: >= 2 failures ■ Evaluation interval: 1 minute **Action File Variables** none Example

Application Health Agent Templates

Application Health agents monitor the new traffic on the network, as well as traffic that is classified into default categories. The specific agents in this category are described in the table below.

> **Note:** You can configure green, yellow, and red thresholds by entering values for just the green and red threshold parameters. The yellow threshold is not a settable parameter–it is the range between the green and red threshold values, and is shown in the table below just for reference.

Agent Template Name	Default Agent Name	Description
Default Traffic	Inbound Default Traffic, Outbound Default Traffic	Agents based on the Default Traffic template monitor the rate (avg-bps) of a Default traffic class, and can alert you when the amount of traffic not classified (falling into Default) is too great. A change in the status of this agent will create a report showing detailed information on the specified class over the most current interval. This agent can only be assigned to default traffic classes, such as */Inbound/Default* or */Outbound/HTTP/Default*. **Agent Parameters** ClassName: Name of the Default class to be monitored RedThreshold: Percent of parent partition's bandwidth that gets classified into the Default class (ClassName); when the rate of ClassName exceeds the RedThreshold, the agent will have a red status GreenThreshold: Percent of parent partition's bandwidth that gets classfiied into ClassName; when the rate of ClassName is less than the GreenThreshold, the agent will have a green status Inbound Default Traffic and Outbound Default Traffic are default agents. Default parameter values are as follows: ■ Red > 15% of parent partition size is classed into ClassName ■ Yellow 7-15% (implied) ■ Green < 7% ■ Evaluation interval: 1 minute **Action Fle Variables** $className, $avg-bps, $class-id Example

Agent Template Name	Default Agent Name	Description
New Application	na	An agent based on the New Application template can alert you when the PacketShaper auto-discovers new applications during the evaluation interval. This agent will return a score value from 0-99999, equal to the number of new auto-discovered classes in the evaluation interval (one day, by default). The red incident report will list the names of all the classes auto-discovered during the interval; a green incident report is not generated. Each unit or PolicyCenter configuration can have only one instance of a New Application agent. **Note:** When discovery is disabled during an interval, the New Application agent's score value will be zero. Therefore, this agent will not work as intended unless traffic discovery is enabled. **Agent Parameters** RedThreshold: Number of new applications discovered; when the number of new classes auto-discovered during the interval exceeds the RedThreshold, the agent will have a red status. GreenThreshold: Number of new applications discovered; when the number of new classes auto-discovered during the interval is less than or equal to the GreenThreshold, the agent will have a green status. Default parameter values are as follows: ■ Red > 0 ■ Green <= 0 ■ Evaluation interval: 1440 minutes **Action File Variables** none Example

Agent Template Name	Default Agent Name	Description
High Bandwidth New App	High Bandwidth New App	An agent based on the High Bandwidth New App template monitors classes that were auto-discovered during the interval, evaluating the amount of bandwidth they are using (avg-bps). This agent can help prevent new applications from consuming too much bandwidth and will create a report when the class crosses a threshold. This agent calculates the amount of bandwidth a newly discovered class is using, as a percentage of the class' parent partition. Typically, the parent partition is either Inbound or Outbound (the link size). But a parent partition isn't restricted to Inbound or Outbound. It can also be a class, such as a subnet class, that has its own partition and has traffic discovery within the class enabled. Each unit or PolicyCenter configuration can have only one instance of a High Bandwidth New App agent. **Note:** When discovery is disabled during an interval, the High Bandwidth New App agent's score value will be zero. Therefore, this agent will not work as intended unless traffic discovery is enabled. **Agent Parameters** RedThreshold: Percent of parent partition's bandwidth; when the rate of any classes auto-discovered during the evaluation interval exceeds the RedThreshold, the agent will have a red status GreenThreshold: Percent of parent partition's bandwidth; when the rate of any newly auto-discovered classes is less than the GreenThreshold, the agent will have a green status The High Bandwidth New App agent is one of the default agents. The default values for the parameters are: ■ Green < 5% of the parent partition ■ Yellow 5-10% of the parent partition (implied) ■ Red > 10% of the parent partition ■ Evaluation interval: 60 minutes **Action File Variables** $namelist, $avg-bps Example

User Event Emulation Agent Templates

There are three types of agent templates in the User Event Emulation agents category: **Class ME Variables**, **Link ME Variables**, and **Partition ME Variables**. These agents monitor a measurement (ME) variable of your choice on a traffic class, link, or partition of your choice. Over 130 measurement variables are available. User Event Emulation agents do not appear on a unit or PolicyCenter configuration by default, but must be manually created. Each User Event Emulation agent is defined by a different set of parameters that determine what is checked, and how the results are reported. The specific parameters for these agents are described in the table below.

Parameter	Value
ClassName	The name of the class, partition or link to be measured (for example, /Inbound or /Outbound/default) Note that *Partition Utilization* agents can only specify the class name of a link or traffic class with a partition; setting a Partition Utilization agent for a class or link without a partition will cause an error. For *Default Traffic* agents, you must specify the class name of default traffic classes, such as /Inbound/Citrix/default or /Outbound/HTTP/default.
MeVariableName	The measurement variable to be monitored. Each measurement type– link, partition, and class– has its own set of variables. The default value is bytes, but any measurement variable of the appropriate type may be used. See a complete list of measurement variables.
MeDuration	The period of time over which measurements will be taken, in minutes.
RedOperator	When you configure a red threshold for a User Event Emulation agent, you must also enter a RedOperator parameter. Use one of the following values: ■ LT– Agent has a red status when the current value of the measurement variable is less than the RedThreshold value ■ LTE–Agent has a red status when the current value of the measurement variable is less than or equal to the RedThreshold value ■ GT– Agent has a red status when the current value of the measurement variable is greater than the RedThreshold value ■ GTE– Agent has a red status when the current value of the measurement variable is greater than or equal to the RedThreshold value ■ EQ– Agent has a red status when the current value of the measurement variable is equal to the RedThreshold value

Parameter	Value
GreenOperator	When you configure a green threshold for a User Event Emulation agent, you must also enter a GreenOperator parameter. Use one of the following values: ■ LT–Agent has a green status when the current value of the measurement variable is less than the GreenThreshold value ■ LTE– Agent has a green status when the current value of the measurement variable is less than or equal to the GreenThreshold value ■ GT– Agent has a green status when the current value of the measurement variable is greater than the GreenThreshold value ■ GTE– Agent has a green status when the current value of the measurement variable is greater than or equal to the GreenThreshold value ■ EQ– Agent has a green status when the current value of the measurement variable is equal to the GreenThreshold value
RedThreshold	The value at which the status of the agent turns red, and a red action file is triggered, if present
GreenThreshold	The value at which the status of the agent turns green, and a green action file is triggered, if present

Access Guided Setup With a Console Connection

If remote-access methods do not work due to network or configuration issues, you can access the unit directly with a null-modem cable. This cable, included with the PacketShaper, offers both 9-pin and 25-pin connectors on each end. To access the command-line interface with a serial connection:

1. Attach the null-modem cable to the serial port on your workstation or PC, using the connector that matches your serial port configuration (9-pin or 25-pin).
2. Connect the 9-pin connector on the other end of the null-modem cable to the PacketShaper s CONSOLE port .
3. Open a terminal emulation program (such as HyperTerminal).
4. Verify that you have configured your program with the following values to communicate with the PacketShaper s console serial port:
   9600 bps, **8** data bits, **1** stop bit, **no** parity, **no** hardware flow control
5. Power on the PacketShaper, if you have not already done so. If the appliance was already turned on, you will need to press **Enter** several times to make the connection. The password prompt appears.
6. Enter the touch user name (**touch**).
7. Enter the password. The default password is **touch**.

   The following message appears:

```
 PacketShaper v11.9 2015-08-06
Copyright (c) 1996-2015 Blue Coat Systems, Inc. All rights reserved.
 PacketShaper not yet configured. Do you wish to be
 guided through initial setup of the PacketShaper? (yes):
```

8. Press **Enter** to start Guided Setup.

Info Tab in the Sky UI

The Info tab provides useful information about your PacketShaper:

- **Access Mode** Touch (read/write) or Look (read only)
- **Alerts** Configuration errors and warning messages for conditions that may affect the operational status of the unit
- **Settings** Status of the PacketShaper's basic settings. To modify any of the settings, you need to go to the Advanced UI – just click the **Change Settings** button to go directly to the Setup tab
- **System Information** Configuration limits, installed keys, memory availability, and time since last reboot (uptime)
- **Product Identification** The product's model number, serial number, management IP address, and the currently running software version

You can use the **Copy to Clipboard** button to copy the information on the Info tab so that you can paste the text into an email message or support ticket.

Policy Inheritance Rules

A traffic class inherits another class' policy if it has no policy of its own.

For each passing traffic flow, PacketShaper traverses the traffic tree looking for a matching traffic class. When found, PacketShaper adjusts the metrics (such as class hits) for the class and applies any associated partition and policies. If the class has no associated policy, PacketShaper continues searching until another matching class is found that does have a policy. The traffic flow *inherits* the second class' policy. But the second class gets no performance metrics recorded other than an additional policy hit.

PacketShaper searches classes in a particular order as it looks for a qualifying inheritable class. It searches *down* the traffic tree, first examining the original class' successive siblings, then the parent's siblings, then the grandparent's siblings, and so on, until a qualifying class is found.

> **Note:** When PacketShaper searches for a class with an inheritable policy, it looks for inheritable classes among sibling classes but not for child classes of those siblings. For example, suppose you have a class, Inbound/Internet/HTTP, with an inheritable policy and another class, Inbound/Subnet_A, that classifies traffic in a subnet and does not have a policy set on it. PacketShaper will classify HTTP traffic in subnet_A into Inbound/Subnet A and then look for an inheritable policy to use. Because PacketShaper will only look at sibling classes (Inbound/Internet in this example, which does not have an inheritable policy), it will not apply the inheritable policy from its child class, Inbound/Internet/HTTP. Thus, to inherit policy for HTTP traffic classified into Inbound/Subnet_A, you can create an Inbound/HTTP class with inheritable policy. Note that you will have to delete Inbound/Internet/HTTP to add Inbound/HTTP or else there will be a conflict and the class will not be evaluated.

PacketShaper searches for a class with four characteristics:

1. Its matching rules encompass the original traffic.
2. It has an associated policy.
3. It is an inheritable class (that is to say, it has an inheritable policy).
4. It is the first such class encountered in PacketShaper's prescribed search order.

For example, given this traffic tree fragment:

- Inbound
  - SAP (with a rate policy)
  - Citrix (with a partition)
  - PeopleSoft over Citrix (no policy)
  - **MS Word over Citrix** (inheritable class with a rate policy)
  - **Default** (inheritable class with a rate policy)
  - SMTP (no policy)
  - **Default** (inheritable class with a priority policy)

/Inbound/Citrix/PeopleSoft traffic does *not* inherit /Inbound/Citrix/MS Word's policy because Word's matching rule conflicts with PeopleSoft's. The /Inbound/Citrix/PeopleSoft traffic *does* inherit /Inbound/Citrix/Default's rate policy. If /Inbound/Citrix/Default were not an inheritable class, then /Inbound/Citrix/PeopleSoft would inherit /Inbound/Default's policy. /Inbound/SMTP traffic inherits /Inbound/Default's priority policy.

> **Note:** Classes that have an inheritable policy have boldfaced names in the traffic tree (as shown in the example above).

In another example, consider the following fragment of a location-based traffic tree with three branch offices and three applications:
- Inbound
- London (Partition of 512 Kbps)
- Paris (Partition of 512 Kbps)
- Rome
- **SAP** (Rate Policy 6)
- **WebBrowsing** (Rate Policy 3)
- **TN3270** (Priority Policy 7)

When a 100 KB TN3270 flow heads to London, the London class gets those bytes added to its usage figures, the flow is paced to fit within London's partition, and the flow gets top access to the partition's capacity with its inherited priority 7 policy. But the TN3270 class' usage is zero because the London class was the initial classification match. The TN3270 class' policy hits does increase by one. You can generate insightful reports about the London class as a whole, but not about the TN3270 class.

To designate a class as inheritable that is, you want other classes to be able to inherit its policy you can either:

- In the Manage tab of the browser interface, select the **Inheritable** button for the class. Note that the class must have a policy before you can see the button.
- In the command-line interface, use the **class set** ***<tclass>*** **inherit** command.

Typically, the default classes such as /Inbound/Default show more policy hits than class hits. That's because the policy hits for the /Inbound/Default class include policy hits for classes without policies that appear earlier in the subtree.

If a class (A) didn't have its own policy and was inheriting a policy from another class (B), and then you assign class A its own policy, there are several policy-related characteristics that will be copied from class B's policy: DiffServ/TOS substitution, delay bound, and flow limit. For example, suppose Inbound/Test was inheriting the Inbound/Default policy (Priority 3, flow limit of 10,000 client and 100,000 server flows per minute) and then you assign Inbound/Test its own policy (a Rate policy). The Inbound/Default class' flow limit policy will be copied to the Inbound/Test class. Thus, Inbound/Test will have a rate policy in addition to a flow limit policy of 10,000 client and 100,000 server flows per minute.

Location-Based Traffic Tree Example

Suppose you want to create location-based classes for a traffic tree at a main site managing traffic for three branch offices at New York, Chicago, and Dallas.

Assumptions:

- Company uses class B IP addresses
- New York and Chicago have T1 WAN links and Dallas has a 512 Kbps line, all full duplex
- All application and Intranet servers are located at the main site

We might have a tree that looks something like the following. The items inside the parentheses are the elements of the class, matching rules, and partition definitions)

Inbound
- New York (Outside, subnet: 128.10.1.0; mask: 255.255.255.0; partition at 1.5 Mbps)
- Chicago (Outside, subnet: 128.10.2.0; mask: 255.255.255.0; partition at 1.5 Mbps)
- Dallas (Outside, subnet: 128.10.3.0; mask: 255.255.255.0; partition at 500 Kbps)
- Default

Outbound
- New York (Outside, subnet: 128.10.1.0; mask: 255.255.255.0; partition at 1.5 Mbps)
- Chicago (Outside, subnet: 128.10.2.0; mask: 255.255.255.0; partition at 1.5 Mbps)
- Dallas (Outside, subnet: 128.10.3.0; mask: 255.255.255.0; partition at 500 Kbps)
- Default

View the Operations Log in the Sky UI

The operations log lists all the command-line interface (CLI) commands that correspond to the actions you perform in the Sky UI. For example, you'll see the commands for setting policies, creating traffic classes, and deleting classes. It's useful to refer to the log if you are having a problem, such as classes not being created.

To view the operations log:

1. Click **Traffic Management**.

2. Click the **Settings** icon. The settings appear in the lower pane, below the class tree.

3. Click the **Operations Log** tab. The CLI commands for the current browser session appear in the log window.

Additional Information

- To clear the current log, click **Clear Log** on the **Operations Log** tab.
- The Policy Manager also offers an **Operations Log** tab, allowing quick, convenient access to CLI commands to troubleshoot as you create classes and policies. This tab has identical log information as the Operations Log in Settings.
- To copy the CLI commands to the clipboard so that you can paste them elsewhere (such as in a support ticket), use the **Copy to Clipboard** button.

Look Up ifName and ifIndex Values

If you don't know the ifName or ifIndex values of your router interfaces, you can query the MIB to determine the values. You'll need an SNMP MIB browser program, such as MG-SOFT's MIB browser.

To obtain the ifName and ifIndex values of a Cisco router interface:

1. Install the MIB files (*.my files downloaded from the Cisco website) that have the definitions of ifName and ifIndex into your MIB browser program.
2. To locate the ifName and ifIndex MIB variables in your MIB tree, walk down the MIB tree, using the object IDs (OIDs) specified below:

   .1.3.6.1.2.1.31.1.1.1.1 (OID for the ifName MIB variable)
   .1.3.6.1.2.1.2.2.1.1 (OID for the ifIndex MIB variable)
3. Run a query of these MIB variables using the query or poll function in your MIB browser program; you should get back the values for these MIB variables.

> **Note:** Some routers may have CLI commands that can determine this information for you, such as the command **show snmp mib ifmib ifindex**, available on some Cisco routers. Check your router documentation for further details.

Matching Rule Details

Overview

A matching rule contains one or more attributes that PacketShaper uses to classify an application's flow. A flow must match all the criteria in the matching rule. Using the **class new** and the **class rule** commands, you can specify a number of arguments that characterize an application's flow.

The matching rule arguments are described in detail in this section.

Protocol

Specify a protocol from the list.

all [|tcp|udp|icmp|ipv6| netbeui|ipx| appletalk|decnet|fna |sna|miscellaneous

Use **all** for non-IP protocols (where inside/outside is not relevant) or for IP protocols if you don't want to match on a specific IP protocol. Use **miscellaneous** or **misc** for other lesser-known protocols, such as RARP and Spanning Tree.

Server Location

Indicate the server location for IP services.

inside *<arg>* IP protocols only.

and/or outside *<arg>*

Use **inside** and **outside** to refer to host locations relative to the unit over the two interfaces labeled INSIDE and OUTSIDE.

You can use a combination of **inside** and **outside** within a single matching rule to match specific hosts from both locations.

Example:

```
class new inbound testboth inside 201.168.0.20 outside 190.30.30.30
```

Hosts

Use an individual specification or a combination of **host**, **service**, **precedence**, and **URL** . If you specify a host without a service, the service defaults to IP that is, match all: TCP, UDP, ICMP, and any port. Likewise, if you specify a service without specifying a host, the host defaults to **any**.

For IP protocols, specify one of:

- dotted-decimal address
- DNS name
- host list relative distinguished name (RDN)

If a DNS server is configured, you can use domain names in place of IP addresses. Also, if you configured a default domain name, you can use a non-fully-qualified domain name if you preface it with the **host:** tag. If the name does not exist when you create the matching rule, PacketShaper warns you about the unknown name. The asterisk wildcard character (*) is supported.

Use one of the following host specifications:

host:any	Type this parameter exactly as shown to indicate traffic from any host.

host:*<ipaddr>*	Use **host:** with an IP address or domain name to classify traffic to or from a specific host.
	
	If you specify a host by domain name, you must either supply a fully qualified domain name or preface the host name with the **host:** tag.
	
	For example:
	`class new inbound test outside tiger.bluecoat.com`
	`class new inbound test outside host:tiger`

net:*<ipaddr>*[:*<submask>*] *<ipv4-addr>*/*<cidr>* or net:*<ipv4-addr>*/*<cidr>*	Specify an IPv4 or IPv6 address or a fully qualified name to specify a single host. Combine an optional subnet mask with an IPv4 address to refer to a set of hosts. .
or net:*<ipv6-prefix>*/*<prefix-length>*	Specify the classless interdomain routing (CIDR) IPv4 address, which enables you to represent an aggregation of networks. Or for an IPv6 subnet, specify the prefix and prefix length. Example: `net:www.mycompany.com/24`
or net:*<dnsname>*/*<cidr>*	
range:*<low>*-*<high>*	Traffic to/from the set of hosts defined by the IPv4 address range of *<low>* to *<high>* inclusive.
local	Indicates traffic to and from (rather than through) the unit itself. Traffic Discovery creates this type of traffic class automatically. **Local** only applies to traffic through the inside interface.
multicast	Match on any IPv4 multicast address. This is a shortcut for specifying a range that matches all class D addresses (224.0.0.0 through 239.255.255.255).

speed: *<low>*-*<high>* Match the detected speed of the host. Specify the host speed as a range of *<low>* to *<high>* in bits per second. To have no upper bound on the speed, specify unlimited as the *<high>* value. This matching rule allows you to allocate different amounts of bandwidth to subpartitions based on the speed of the client, or assign different policies depending on whether you have high-speed or low-speed clients.

Examples:

```
class new /outbound slow outside speed:0-56k IP
```

```
class new /inbound newsbatch inside speed:200k-unlimited NNTP outside
speed:200k-unlimited
```

Additional Information

- The browser interface does not currently offer a way to create a host speed matching rule. Therefore, if you create a class in the CLI using the speed specification, the browser interface has no way of showing this rule; instead, it shows as a match-all class. In addition, modifying the matching rule in the browser interface will remove the host speed specification.
- There can be a delay between the time a host is detected on the network and the computation of its speed capabilities, resulting in temporary misclassification.

mac: *<macaddress>* Enter a MAC address for IP or non-IP protocols to enable classification of traffic to and from a known host. Specify a MAC address in six-octet format – for example, *08:12:34:56:A1:5C*. Note that you cannot classify Ethernet broadcasts (*FF:FF:FF:FF:FF:FF*).

Examples:

```
class new inbound server1 all mac:08:12:34:56:A1:5C
```

```
class new inbound server2 all mac:01:23:45:67:89:AB
```

Note: The **all** parameter classifies both IP and non-IP protocols. Although it is possible to specify inside or outside matching rules, keep in mind that classes created with those rules will catch only IP-based traffic. For example:

```
class new inbound server3 inside mac:01:23:45:67:89:DE
```

In watch mode, you need to create two matching rules to classify by MAC. Even if you know the MAC resides on the outside, you need one matching rule with the MAC address on the inside and the other matching rule with the MAC address on the outside. For example, to create an Inbound class named mac_router:

```
class new inbound mac_router nodefault all mac:inside:01:23:45:67:89:DE
 class rule add inbound/mac_router all mac:outside:01:23:45:67:89:DE
```

When the PacketShaper is inline, you need only one matching rule.

list: *<name>* Specify a host list name. To create a host list, use the hl new command.

Service Groups

A *service group* is a set of services that have been grouped together based on a common functionality. For example, all services that classify online games are in the Games group. Service groups offer a quick way to classify and control a group of related services. For example, suppose you want to limit peer-to-peer traffic on your network. Previously, you would have had to create

classes for each P2P service and apply the same policy to each class. With the service group feature, you can create a class based on the P2P service group and apply a policy to this single class.

group: *<group name>*	Substitute a built-in or custom group name for *<group name>*. Use the class group show command to list valid service group names. Example: `class new inbound games inside group:games outside group:games`

Services

Specify a service in a number of ways, as shown in the following list. Applications are assigned port numbers by the Internet Assigned Numbers Authority (IANA). For details about port number assignments, refer to RFC 1700, *Assigned Numbers*. Note that services are server based, so you must consider where the server of the service is located (inside or outside).

service:any	Match traffic using any protocol TCP, UDP, or ICMP and any port.
service: *<item>* or *<item>*	Substitute a service name for *<item>*. The service name can be one of the built-in services or one of WebPulse's web application services. To see a list of all services (built-in and WebPulse), use **class services**. To see a list of web application services only, use **setup webpulse show applications**. Example: `service:smtp`
operation: *<operation name>*	Classify by a specific web operation. (Note: Requires WebPulse to be enabled.) Can be used in conjunction with a service to classify traffic for a web operation within a web application. For example, the following matching rule classifies message posting in LinkedIn: `class new inbound linkedin_messages outside service:LinkedIn operation:Post_Messages`
proxy: *<service>* port: *<number>*	Identify an application that is running on a nonstandard port. Example: `class new /outbound/outside news outside host:news-server.com proxy:NNTP port:11199`
tcp:any or tcp or udp:any or udp	Match all TCP or UDP traffic from any port.

URL Categories

Web URLs can be grouped into various categories, such as social networking, gambling, pornography, news media, and shopping. PacketShaper is able to analyze URLs that users are requesting, determine what category the website belongs to, and classify the traffic into the appropriate category class. This gives you granular visibility into the type of web traffic on your

network. With the URL category feature, you can create a class based on a specific category and apply a policy to this single class.

urlcategory: *<category_name>*|all

Enter a specific *<category_name>* or specify **all** to create a class that matches all categories. Use the setup webpulse show categories command to see a list of valid category names. If the category name has a space, enclose the name in quotes.

Example:

```
class new outbound web_apps inside urlcategory:”web applications”
```

Users

When PacketShaper is configured as a BCAAA client, you can manually create a class for a specific Active Directory user or group, or for a user list that has been defined in PacketShaper.

user-group: *<group name>*

Specify a user group name in the *<domain-name>**<group-name>* format. You must enter the domain name; it is not optional. You may enter UTF-8 characters for the group name (but not for the class name).

Example:

```
class new inbound FinanceGroup inside user-group:finance\group-finance
```

Note: Each PacketShaper model has a limit to the number of group-based classes that can be created. Once you have exceeded the maximum number of user groups that can be used in matching rules, you must delete an existing group-based class before you can create a new one.

user-name: *<user name>*

Specify a user name in the *<domain-name>**<user-name>* format. Use the class users command to see a list of valid user names. If the user name has a space, enclose the name in quotes. You must enter the domain name; it is not optional. You may enter UTF-8 characters for the user name (but not for the class name).

Example:

```
class new outbound john.smith inside user-name:"finance\john smith"
```

user-list: *<user list>*

Specify the name of a user list. Use the ul show command to see a list of defined user lists.

Example:

```
class new inbound myuserclass inside user-list:myuserlist
```

Specific Sub-Protocol

This option is available only via the command line and traffic discovery.

protocol: *<number>*

Specify a particular subprotocol of IP, DECnet, or Miscellaneous. When you use this protocol parameter, do not specify a host, port, or address.

If the protocol number is greater than 256, command output (such as class show) will display a 4-digit hex number.

Example:
`protocol:0x6003`

For protocol numbers up to 256, a decimal number is shown in the class show output.

Example:
`protocol:233`

Web-Specific Parameters

To identify specific web traffic, use the following parameters:

http: Specify web traffic. Web traffic is defined as: TCP protocol and HTTP service.

web:url:"*<url>*" Classify traffic for a specific URL, listed within quotes. The URL must be specified on the server side. The wildcard characters (***** for a group of characters and **?** for single characters) are supported. Slashes are not treated as special characters. To match a specific URL (without wildcards), the URL in quotes must be the full pathname, including the leading slash. Only one URL can be specified within a matching rule.

Example:

To create a class that goes to the URL http://altman.com/support/support.htm, use this command:

`class new inbound altman outside service:http host:altman.com web:url:"/support/support.htm"`

web:content-type: Classify traffic according to the content-type string in the HTTP get or put header.

Example:
`web:content-type:"image/gif"`

web:user-agent: Classify traffic according to the requesting process, possibly a web browser, or robot.

Example:
`web:user-agent:"Mozilla 4.0"`

web:host: Classify traffic according to a specific DNS name or IPv4 address. The wildcard characters (***** for a group of characters and **?** for single characters) are supported.

Example:
`web:host:192.68.254.12`

Application-Specific Criteria

The following services can be further classified by specific criteria: Citrix-ICA, DCOM, DICOM, FTP-Data-Clear, FTP-Data-Clear-IPv6, HTTP, HTTP-IPv6, HTTP-Tunnel, HTTP-Tunnel-IPv6, ICMP, ICMPv6, NNTP-Clear, Oracle-netv2, PostgreSQL, ProxySG-ADN, ProxySG-ADN-IPv6, RTCP-I, RTP-I, SMTP-Clear, SMTP-Clear-IPv6, SOAP-HTTP, SSL, SSL-IPv6, and WAP .

<application>:*<attr>*:*<value>*

Classify according to an attribute specific to the application. *<value>* cannot exceed 99 characters. Use the class criteria attributes command to list valid application and attribute names. Use the class criteria track and class criteria recent commands to determine the *<value>*.

Note: You must specify the service type when specifying application-specific criteria.

Examples:

```
class new inbound/citrix peoplesoft inside service:citrix-ica
citrix:application:peoplesoft

class new inbound/http web-graphics inside service:http web:content-
type:image/gif

class new inbound/ftp ftp_mp3_downloads outside service:ftp-data-clear
ftp:filename:*.mp3
 class rule add inbound/ftp/ftp_mp3_downloads inside service:ftp-data-clear
ftp:filename:*.mp3
```

Device

To match traffic that is passing through a specific interface, you can create a matching rule that specifies a physical port on the PacketShaper. For example, you can create a class that classifies all outbound traffic on the NIC in slot 2. You can then turn on class discovery for the NIC class and PacketShaper will auto-discover all the applications, services, and protocols on this NIC.

device: *<device>*

Classify according to a network interface. Valid values for *<device>* depend on your PacketShaper model. For example:

slot1
slot2
slot3
slot4

If four-port NICs are installed, you designate a port pair (pair1 or pair2). For example:

slot3_pair1
slot3_pair2
slot4_pair1
slot4_pair2

Examples:

```
class new inbound slot2_NIC all device:slot2

class new outbound slot3_left_NIC all device:slot3_pair1
```

IP Precedence (Diffserv)

The following parameters can be defined in a matching rule to match traffic with IP precedence values set.

dscp	IP protocols only dscp:*<value>*|any Enter an optional Differentiated Services code point (DSCP) value (0-63), a range of values, or the literal **any** . The DSCP is the first six bits in the TOS byte in the IP header.
precedence: *<value>* |*<range>*|any	IP protocols only Enter an optional precedence value (0-7, where 7 is the highest priority), a range of values, or the literal **any**.
tos:*<value>*|any	IP protocols only Enter a value (0-15) or the literal any. Since the standard for this field is changing, it is recommended that you choose one number from the following list, rather than combining values. 8 = minimize delay 4 = maximize throughput 2 = maximize reliability 1 = minimize monetary cost 0 = normal service

Virtual LAN Identification (802.1Q and ISL)

PacketShaper has the ability to classify application traffic within multiple virtual LANs on a single Ethernet segment (an 802.1Q trunk) or ISL (Inter-Switch Link, a Cisco protocol). For a multiple VLAN environment (where packets passing through the unit will have more than one VLAN header), only the outermost VLAN header is used for classification. The other headers are ignored. If packets will be tagged with multiple VLAN headers, be sure to verify that the additional header information does not create a final packet size greater than 1548 bytes before the packet reaches the PacketShaper. (PacketShaper can accept packets up to 1548 bytes in size and will drop any packets that exceed this size.)

```
[vlantype:ISL|8021Q]
 [vlanid:<value>|<range>|any]
```

Indicate the Virtual LAN (VLAN) identification value (or range of values) of the traffic you want to match. Valid **vlanid** values for 802.1Q are 0 - 4095. For ISL, valid values are 0 32767.

For example, to create a new class for all VLAN 802.1Q traffic with an ID tag of 24:

```
class new inbound VLAN24 all vlantype:8021Q vlanid:24
```

Once the class is created, you can enable discovery within the class to discover all applications running over this VLAN.

> **Note:** When creating a VLAN class in the CLI, you cannot specify all three VLAN matching rules (vlantype, vlanid, and vlanpriority). If you want to specify the 802.1Q VLAN type, you can simply leave out the **vlantype:8021q** rule because it's the default value when a type is not specified. Alternatively, you can create the class in the browser interface.

VLAN Priority (802.1p)

To classify according to VLAN priority (802.1p) , use the following parameter when creating a class or adding a matching rule:

`[vlanpriority:<value>|<range>|any]`

Indicate the priority value (or range of values) of the traffic you want to match. Valid **vlanpriority** values for 802.1p are 0 7.

For example, to create a new class for all VLAN 802.1p traffic with a priority of 5:

```
class new inbound priority5class all vlanpriority:5
```

Modify a Host List

After you have created a host list, you can modify it add members, delete members, delete host lists, or copy host lists.

To modify a host list:

1. Click **edit list** in the *New Traffic Class* window, the *New Matching Rule* window, or the *Edit Matching Rule* window. The *Host List Editor* window opens. show screen

   - To add members to a host list, select the host list name, enter an IP address, subnet, IP address or subnet range, or DNS name in the **New** hosts field, and click **add host**.
   - To delete a member from a host list, select the item and click **delete host**.
   - To delete a host list entirely, select the host list name and click **delete list**.
   - To copy a host list, select a host list, supply a new name, and click **add list**.

2. Click **done** to return to the *Matching Rule* or *Traffic Class* window.

3. Click **apply changes**.

NetFlow v5 Header

The following fields are in the UDP header for NetFlow v5:

Name	Description	Offset	Field Length in Bytes
Version	Type of record format (5=NetFlow-5)	0	2
Count	Number of flow records contained in the UDP packet (maximum of 30)	2	2
SysUpTime	Value of SysUpTime) when the UDP packet was sent (measured in milliseconds); SysUpTime measures how long the system has been up and running	4	4
Epoch Seconds	Number of seconds between the epoch (January 1, 1970) and the time when the UDP packet was sent. (Epoch is a date used as the "beginning of time" for timestamps. Time values in Unix systems are represented as the number of seconds since the epoch.)	8	4
Nanoseconds	Residual nanoseconds (billionth of a second) after the epoch second	12	4
Flows Seen	Number of flows seen since PacketShaper began emitting flow detail records	16	4
Engine Type	User-configurable value (0-255) assigned to the PacketShaper that is emitting flow detail records. Use the setup flowrecords engineType command to assign a type to the unit.	20	1
Engine ID	User-configurable value (0-255) assigned to the PacketShaper that is emitting flow detail records. Use the setup flowrecords engineID command to assign an ID to the unit.	21	1
Sampling Info	n/a	22	2

NetFlow v5 Record Format

The following fields are recorded in the NetFlow-5 record type:

Name	Description	Offset	Field Length in Bytes
Source IPaddr	IP address of the device that sent the flow	0	4
Destination IPaddr	IP address of the destination device	4	4
Next hop router IP address	n/a	8	4
Inbound snmpIFindex	SNMP index number that identifies the Inbound interface on the PacketShaper.	12	2
Outbound snmpIFindex	SNMP index number that identifies the Outbound interface on the PacketShaper.	14	2
Packet Count	Number of packets in the flow	16	4
Byte Count	Total number of bytes in the flow	20	4
Time at Start of Flow	Value of SysUpTime when the first packet in the flow was seen (measured in milliseconds)	24	4
Time at End of Flow	Value of SysUpTime when the last packet in the flow was seen (measured in milliseconds)	28	4
Source Port	Port number of the device that the flow went out of	32	2
Destination Port	Port number of the device that the flow went to	34	2
One pad byte	n/a	36	1
TCP flags	Protocol state (URG=32, ACK=16, PSH=8, RST=4, SYN=2, FIN=1). For example, a value of 27 indicates the flow had a SYN, ACK, PUSH, and FIN (2+16+8+1=27).	37	1
Layer 4 Protocol	Type of layer 4 protocol. For example, ICMP=1, TCP=6, Telnet=14, UDP=17	38	1
IP Type of Service (ToS) / Diffserv	Value that designates special handling of traffic (precedence, delay, throughput, and reliability)	39	1
Source Autonomous Sys ID	n/a	40	2

Name	Description	Offset	Field Length in Bytes
Dest. Autonomous Sys ID	n/a	42	2
Source Mask Bits Count	n/a	44	1
Destination Mask Bits Count	n/a	45	1
Two Pad Bytes	n/a	46	2

Certain fields (those marked with *n/a* in the Description column above) are applicable to routers, but not PacketShapers. These fields will always have a value of zero (0) in the NetFlow v5 records from the PacketShaper.

PacketShaper Folders

The PacketShaper's public/9.256 folder contains the following directories:

Folder Name	Contents
CFG	Configuration files. For example, config.ldi contains the traffic tree configuration (including all classes, class IDs, partitions, policies, host lists, and events) and settings.cfg contains unit-specific settings (such as IP address, DNS server, NIC information, and domain name).
LOG	Log files that are automatically generated when the unit boots or crashes; used by Customer Support for diagnosing problems
CMD	Command files– either user-created or generated with the setup capture command. See also Create a Command File.

The public/9.258 folder contains these folders:

Folder Name	Contents
PKTLOG	TCPDump files (.DMP) created by the packet logging feature. See Sniff Without a Sniffer.
DIAG	Diagnostic information that is periodically gathered about the unit; used by Customer Support for diagnosing errors and problems
ME/MEASURE	Measurement data files for links, partitions, and classes (daily and monthly)
CMD	Adaptive response action file output files
IMAGE	PacketShaper image files
LOG	Log files created by the user event feature. See View Event Logs.
AGENT	Adaptive response agents, agent templates, and action files.
SUBSCRIBE	When the unit is subscribed to PolicyCenter, this folder will contain the images that are distributed to the unit.
TMP	Temporary files

The public/9.1026 folder contains these folders:

Folder Name	Contents
URLCAT	The URL category cache and category map. See URL Categories Overview.
WEBAPP	XML files associated with the web application classification feature.

The public/3.0 folder has an SSL folder where digital certificates are securely stored. Instead of using PacketShaper's self-signed certificate, you can generate your own PEM file using OpenSSL library, copy it to the SSL folder, and then import it with the "setup ssl install" on page 1345 CLI command.

To copy files to and from the PacketShaper, use the File Browser.

PacketShaper SNMP Traps

PacketShaper supports up to eight destinations for SNMP traps. Configure these addresses via the browser (see Configure PacketShaper for SNMPv1-v2c Support or Configure PacketShaper for SNMPv3 Support) or the command-line interface.

PacketShaper issues several traps, which provide notification of component failures or status changes for certain PacketShaper models.

BCAAA Primary Status A trap is generated when the primary BCAAA server connected to PacketShaper changes to or from a state of less readiness.

- OK (1) – The status returned to normal after a failure condition.
- Fail (2) – PacketShaper cannot log in to the primary BCAAA server.
- Unconfig (3) – A primary BCAAA server is not configured.

> **Note:** PacketShaper uses the Symantec Authentication and Authorization Agent (BCAAA) to convert IP addresses to user names so that it can classify and report on network users.

BCAAA Secondary Status. A trap is generated when the secondary BCAAA server connected to PacketShaper changes to or from a state of less readiness.

- OK (1) – The status returned to normal after a failure condition.
- Fail (2) – PacketShaper cannot log in to the secondary BCAAA server.
- Unconfig (3) – A secondary BCAAA server is not configured.

Fan Failure A trap is generated when any of the fans fail. The trap indicates the number of the fan that failed.

High Availability Device Down. A trap is generated when a device that is being monitored by the high availability feature goes down. The device can be an access router or a WAN link interface. The trap includes the following information:

- Status – 0 (meaning the device has gone down)
- Router Name – The IP address or DNS name of the configured router that has gone down or whose configured interface has gone down
- Interface Name – The name of the interface (ifname) that has gone down. If the interface name could not be resolved, UNKNOWN will be displayed. If the router went down (as opposed to an individual interface), no information will be displayed.

License Expired PacketShaper sends a trap after a Symantec license has expired, and continues sending traps once an hour until the license is updated.

Measurement Engine A trap is generated when PacketShaper detects any of the following conditions:

- The measurement engine fails to start properly.
- Startup is disabled.

- Measurement data cannot be saved because of a problem on the hard drive.

Power Supply Failure A trap is generated when a power supply fails. The trap OID corresponds to the power supply number that has failed.

Site Router Status. A trap is generated to indicate the status of PacketShaper's site router:

- ready – Site router is actively routing or ready to route.
- gone – PacketShaper is not receiving ARP replies from the router.
- error inside – PacketShaper received an ARP reply, but from the inside.
- none – No site router has been configured.

Temperature PacketShaper has sensors for monitoring the temperature of the CPU, RAM, motherboard, and power supplies. If a component's temperature exceeds the proper range, an SNMP trap is sent. Example message: *badTempMbAndRamAndCpu(7)*

Voltage Status The 1V, 3V, 5V, 12V, and CPU sensors are monitored. Each platform typically has multiple sensors of each type (for example, there are multiple 1V sensors). If one or more voltage regulators falls outside of the proper range, an SNMP trap is sent. Example message: *badVoltage12v(4) The 12v input is abnormal*

In addition, all of the MIBs supported by Net-SNMP 5.5 are available, including the following standard SNMPv2 traps:

Cold/Warm Start PacketShaper issues a MIB-II trap during reset– that is, when loading a new image, manually initiating a unit reset, or recovering from a crash. A *cold start* trap is sent when the PacketShaper is power-cycled or is without power for any amount of time. A *warm start* trap is sent for software resets.

Link Up/Down A MIB-II trap is generated each time a link goes down and again when the link returns to normal. The trap indicates which link went down.

> **Tip:** If you aren't interested in receiving notifications about a particular SNMP trap and want to prevent PacketShaper from sending the trap, you can set up a trap filter. See "SNMP Trap Filtering" on page 1133 for more information.

Arithmetic Comparisons

Arithmetic comparison is required for certain portal features for example, display a red symbol if one of the numbers exceeds a threshold.

The > and < symbols are used by HTML, so the portal does not use those symbols in tags. Instead, it is more like Fortran:

Operator	Description
.LT.	Less than
.LE.	Less than or equal to
.EQ.	Numeric equals, not necessarily the same as string equals. See below for number syntax.
.NE.	Numeric not equals
.GT.	Greater than
.GE.	Greater than or equal to

Numeric constants are somewhat more flexible than just numbers: a rate syntax is supported as well. Entries such as "2M" and "1.24K" are accepted. The percent sign (%) is also allowed, so that comparisons on tcp-efficiency% and other percentage-based measurement engine variables can be allowed. Constants must be quoted if they contain any non-numeric characters such as "M".

Arithmetic Example

A single table row from the sample EX1.HTM file is shown below. In this example, IF statements are combined with arithmetic comparisons to display a red, yellow or green table cell, based on the inbound throughput.

```
<TR>
<TD> Inbound Throughput (bps) </TD>
<TD align=right BGCOLOR=
%IF (%PCLASSLIST.NAMED[0.AVG-BPS.LAST-MINUTE]% .GT. "2M")%
"red"
%ELSEIF (%PCLASSLIST.NAMED[0.AVG-BPS.LAST-MINUTE]% .GT. "1.9M")%
"yellow"
%ELSE%
"green"
%ENDIF%>
 
</TD>
</TR>
```

Select Traffic Classes for the Portal Page

The service provider authoring a portal page for a particular customer must know which traffic classes are of interest to that customer. This is one of the ways that selective disclosure is implemented. Up to 20 classes can be listed on a portal page.

The %PUSH% statement with the PCLASSLIST object defines the traffic classes that are to be reported on the portal page. It must appear at the beginning of the portal page.

The PCLASSLIST (for portal class list) is used as follows:

```
%PUSH (PCLASSLIST="/Inbound;/Outbound;/Inbound/HTTP;/Outbound/HTTP")%
```

You can list up to 20 traffic classes in the %PUSH% statement; each class name must be separated by a semi-colon, and quote marks must surround the list of class names.

The PCLASSLIST object is the only tag language context needed to build the portal pages, so it's worth explaining exactly what it does:

- For a given set of traffic classes it fetches in one pass all the measurement engine variables needed and all their date functions. Since the measurement engine data set is fairly extensive, this optimizes the performance.
- It computes a set of date functions on each variable, such as "last minute,""last 60 minutes," "yesterday," and "this month."
- It provides a set of tags to allow outputting any of this information to the HTML page.

There is one PCLASSLIST object per page and the first appearance of it triggers a fetch from the measurement engine of all the data required. Subsequently there can be many more uses of the PCLASSLIST object for example, %LOOP% statements to output multiple tables, without incurring any more measurement engine overhead.

Search Facility

You can use a search command while fetching the list of traffic classes for a page. This allows you to have a single portal page for multiple customers and/or multiple PacketShapers. The OWNER1.HTM file that you downloaded illustrates a page that can be used for multiple customers.

The search command can be used as follows:

```
%PUSH (PCLASSLIST="SEARCH:OWNER=<search string>")%
%PUSH (PCLASSLIST="SEARCH:NAME=<search string>")%
%PUSH (PCLASSLIST="SEARCH:NOTES=<search string>")%
```

The owner and notes (comments) are specified for a class in the **Manage** tab in the Advanced UI or with the **class owner** and **class note** commands in the CLI. *<search string>* may either be a literal, or the supported tag language construct:

```
%CUSTOMER.NAME%
```

`%CUSTOMER.NAME%` evaluates to the name of the logged-in portal user.

For each of the three search options (Owner, Name, and Notes), the list of traffic classes is those for which the given string appears anywhere in the relevant field, case-insensitive. For example:

Name	Owner	Notes
/Inbound		Never touch!
/Inbound/HTTP	Jo	
/Outbound	Joseph	Really never touch!
/Outbound/HTTP	Mary JO	Mary JO created this

Search strings behave as if bracketed by wild cards. For example, these are the results of various queries on the traffic tree shown above:

Query	Results
OWNER=JO	/Inbound, /Outbound, /Outbound/HTTP
NOTES=never	/Inbound, /Outbound
NAME=HTTP	/Inbound/HTTP, /Outbound/HTTP
OWNER=Mary JO	/Outbound/HTTP
NOTES=touch	/Inbound, /Outbound
NOTES=mary	/Outbound/HTTP

Using the Same Portal Page on Multiple Units

To use the same portal page on different units, create a customer portal page that includes a statement similar to the following:

```
%PUSH (PCLASSLIST="SEARCH:NOTES=PortalPage")%
```

This generic portal page can be copied to each applicable unit. For this technique to work, you add the text *PortalPage* in the Comment field of each class that you want to appear on the portal page for that unit. When the customer accesses the portal page, he will see a customized page that includes a list of classes that have the *PortalPage* comment.

Using Different Portal Pages on a Unit

The customer portal feature can also accommodate multiple portal pages on the same unit, in the case where multiple customers are sharing a single PacketShaper and you want each customer to have a different portal page. For example, customer A's portal page contains the following statement:

```
%PUSH (PCLASSLIST="SEARCH:NOTES=CustomerA")%
```

and customer B's page contains this statement:

```
%PUSH (PCLASSLIST="SEARCH:NOTES=CustomerB")%
```

The portal pages can be copied to the appropriate customer's directory in the 9.258/customer directory. Then, you add the text *CustomerA* to the Comment field of each class that is applicable to customer A, and add the *CustomerB* comment for classes applicable to customer B. If a class is applicable to both customers, enter *CustomerA and CustomerB* for the comment. When customer A accesses the customer portal, his page will list all classes with *CustomerA* as a comment. When customer B accesses the portal, he will see a list of all classes with *CustomerB* as a comment.

Display Custom Information on the Portal Page

To insert the customer's name in a page, use:

```
%PCLASSLIST.CUSTOMER.NAME%
```

To insert the message of the day anywhere on the page, use:

```
%PCLASSLIST.CUSTOMER.MOTD%
```

This inserts the MOTD for the currently logged-in customer. If the customer is assigned a specific Message-of-the-Day, this string will return their custom message. If a customer account is configured to use the global Message-of-the-Day, it will be returned.

To insert the earliest and latest measurement engine data that was found, use:

```
%PCLASSLIST.TIMESTAMP.START%
```

and

```
%PCLASSLIST.TIMESTAMP.END%
```

Note that these tags will work only if the page has the %PUSH% statement with the PCLASSLIST object. This statement defines the traffic classes that are to be reported on the portal page, and must appear at the beginning of the portal page. If you don't want to list any classes on the portal page, you can specify an empty list. For example:

```
<html>
 <body>
 %PUSH (PCLASSLIST="")%
 %PCLASSLIST.CUSTOMER.NAME%
 </body>
 </html>
```

The above code displays the name of the logged-in portal user on the portal page.

Include Graphs on a Portal Page

The *OP.MEAS.GETGRAPH* command is the basic means of incorporating PacketShaper graphs into the report.

Each graph has its own URL. Here is an example of how a graph could be embedded in the HTML pages:

```
< IMG SRC=/customer/Cgi1?OP.MEAS.GETGRAPH=&MEAS.NAME=/Inbound&MEAS.TYPE=link&
MEAS.PERIOD=3600&MEAS.VARLIST=avg-bps%%20
peak-bps&MEAS.VARLISTLABEL=AverageRate%%2CPeakRate&MEAS.GRAPHTYPE=line >
```

> **Tip:** In the code above, %20 is used to represent the space character, and %2C is used for a comma. If spaces or commas are in text that is enclosed in quotes, it is not necessary to represent them using %20 or %2C.

The /customer/Cgi prefix is mandatory for these graphs; it tells PacketShaper to allow the *OP.MEAS.GETGRAPH* operation for authenticated customers.

In this example, a graph is included with the following characteristics:

CGI Parameter	Value (in above example)
MEAS.NAME	/Inbound
MEAS.TYPE	Link (as opposed to partition or class)
MEAS.PERIOD	A period of 3600 seconds (1 hour)
MEAS.VARLIST	Two ME variables, avg-bps and peak-bps. Note that they are separated by a URL-encoded space, %20, and the % must be doubled (because % characters are normally interpreted by the tag engine).
MEAS.VARLISTLABEL	Two labels for the two ME variables. They are separated by a URL-encoded comma character, %2C. Again, the % must be doubled.
MEAS.GRAPHTYPE	A line (as opposed to bar or pie) graph

The maximum number of classes that can be specified when creating a stacked chart for a customer portal is 16.

Saving the Web Page

These web pages can be saved using your browser's web page complete option, unless the author uses Javascript's "document.write" function to create dynamic HTML.

If the author avoids JavaScript and follows a few simple rules, a report with embedded graphs can be saved, with all its graphs.

Here is an example of a report with two embedded graphs, after being saved in Internet Explorer:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.0 Transitional//EN">
<!-- saved from url=(0042)http://10.7.2.0/customer/mycust/glast60.htm -->
<?xml version="1.0" encoding="UTF-8"?><HTML lang=en xml:lang="en"
xmlns="http://www.w3.org/1999/xhtml"><HEAD><TITLE>Graph sample</TITLE>
<META http-equiv=Content-Type content="text/html; charset=windows-1252">
<META content="MSHTML 5.50.4522.1800" name=GENERATOR></HEAD>
<BODY>
```

```
<H1>Network performance overall, last 60 minutes </H1>
<H2>Inbound </H2>
<TABLE>
<TBODY>
<TR><IMG src="Graph sample2_files/Cgi1.gif">
<IMG src="Graph sample2_files/Cgi2.gif"> </TR></TBODY></TABLE></BODY></HTML>
```

Note that the two graphs were saved in a subdirectory, Graph sample2_files, as Cgi1.gif and Cgi2.gif, and the URLs changed to link to those files instead of linking back to the server. This is the special significance of the 1 in Cgi1 above: IE uses it as a file name. In this two-file example, the two URLs started with /customer/Cgi1 and /customer/Cgi2.

When you open this HTML file with IE, it's displayed without going back to PacketShaper. You can email it after using a file compression utility such as WinZip to put all the files into a recursive Zip file.

Link to Other HTML Pages

The HTML pages in customer home directories may contain links to a very limited set of information within PacketShaper–those files that have been placed in the home directory, any public files, and a CGI which generates graph images. In addition they can contain links to anywhere outside PacketShaper– to the service provider's own website, for example.

A public file is any file placed in the 9.258/customer directory. A typical use of this would be the GIF and JPEG files for the service provider's banner, or general information that applies to all customers.

Public files can be referenced with a relative URL. An example for the portal might be:

```
<P>Please follow these <A href="../problem.htm"> procedures </A> to report a problem </P>
```

Display Statistics on the Portal Page

Error Detection

It is important to test the status of the data to be fetched before you start. Inevitably, things can go wrong with fetching the data from the measurement engine. The tags might have specified a traffic class that no longer exists on the device, or never did; the measurement engine might not have started yet; disk errors or other problems might exist. Therefore, PacketWise provides a way for the service provider to test the results of the measurement engine fetch. If it was not successful, you can output an error message or cut the page short.

%PCLASSLIST.ERRORSTATUS% gives the measurement engine's results. Here is an example of how a service provider might handle it in the HTML code:

```
 %IF (%PCLASSLIST.ERRORSTATUS% != "")%
<h3><b><div align="left"><font color="Red">
Errors encountered while preparing this report:<br>
 %PCLASSLIST.ERRORSTATUS% <Br>
The rest of your report will be rendered as well as is possible, but there may be missing sections.
Please contact Technical Support with this problem.
</font>
</div></b></h3>
%ENDIF%
```

You also could enclose the rest of the report within an %ELSE% statement, so that it only comes out if there are no errors. The portal attempts to recover from errors and deliver some data if it can, but there are some errors from which there is no real recovery.

Measurement Engine Data

This section explains how you can take the data that the portal has gathered on classes, and insert it into the HTML page.

> **Note:** Due to tag engine restrictions, the classes are not actually named but numbered.

Data from the *PCLASSLIST* context is referenced in one of two ways:

- To retrieve data from a named class such as /Inbound, /Outbound, or /Inbound/HTTP:
  %PCLASSLIST.NAMED[*<named-class-number>*.*<me-var>*.*<date-function>*]%

- In a *%LOOP%* statement, to refer to the current entry:
  %PCLASSLIST.CURRENT[*<me-var>*.*<date-function>*]%

where

<named-class>—a number from 0 to the number of classes minus 1

<me-var>—measurement engine variables

<date-function>—one of the following date functions:

Function	Description	
LAST-MINUTE	Previous minute	one minute
LAST-15-MINUTES	Previous 15 minutes	one minute
LAST-60-MINUTES	Previous 60 minutes (for example, if the current time is 11:15, the LAST-60-MINUTES would be 10:15-11:15)	one minute
LAST-FULL-HOUR	Previous full hour (for example, if the current time is 11:15, the LAST-FULL-HOUR would be 10:00-11:00)	one hour
LAST-24-HOURS	From 24 hours ago to now	one hour
LAST-7-DAYS	From seven days ago to now	one hour
TODAY	From midnight to now	one hour
YESTERDAY	Previous day, from midnight to midnight	one hour
THIS-MONTH	From midnight on the first day of the current month to now	one hour
LAST-MONTH	Previous full month	one hour

Examples:

%PCLASSLIST.NAMED[0.AVG-BPS.LAST-MINUTE]% —returns the value of the avg-bps variable for the last minute, for the /Inbound traffic class (which was the first class in the list, thus the 0).

%PCLASSLIST.NAMED[1.TCP-EFFICIENCY-PCT.LAST-MONTH]% —returns the tcp-efficiency% variable for /Outbound (the second class in the list), for the last month.

%PCLASSLIST.CURRENT[SERVER-DELAY-AVG.LAST-60-MINUTES]% —within a %LOOP%, returns the value of server-delay-avg for the current class, for the last 60 minutes.

%PCLASSLIST.NAMED[0.AVG-BPS.TODAY]%—returns the value of the avg-bps variable for the first class in your list, since midnight of the current day.

Displaying Current Class Name

This example illustrates a method for displaying the current class name, within a table, using a loop structure to fill the table with a row for each class:

```
<TABLE>
<TR>
<TH>Class</TH>
<TH>Bytes</TH>
</TR>
%LOOP(PCLASSLIST="...listofclasses...")%
```

```
<TR>
<TH>%PCLASSLIST.CURRENT[NAME]%</TH>
<TD>%PCLASSLIST.CURRENT[BYTES.LAST-MINUTE]%</TD>
</TR>
%ENDLOOP%
</TABLE>
```

In the example above, the list of classes is used in the %LOOP% statement, and for each class in the list, a new row is generated in the table. %ENDLOOP% is inserted before the closing </TABLE> tag.

Meta-data

In addition, there can be an optional last parameter within the [ ] pair, the meta-data tag. Meta-data is data about data. The available meta-data tags are:

VALUE—the default; no operation

SAMPLE-COUNT—the number of measurement engine samples that fell into this bucket. (Useful if, for example, you had the unit on for only half of last month, so you aren't seeing a full month's worth of data)

EARLIEST-TIME—the earliest time stamp on the samples

LATEST-TIME—the latest time stamp on the samples

TIME-SPAN—the time, in seconds, between the earliest and latest sample found

Examples:

%PCLASSLIST.NAMED[0.INBOUND.AVG-BPS.LAST-MONTH.LATEST-TIME]%— returns the latest sample found as a date-time string.

%PCLASSLIST.NAMED[0.AVG-BPS.TODAY.TIME-SPAN]%—returns the time, in seconds, between the earliest and the latest ME samples that were used to compute AVG-BPS.TODAY for the first class in the list.

> **Tip:** Meta-data may be used for a page with a link labeled "about this data" which leads to an appendix that shows how the data was derived.

Customer Portal API

In addition to any valid HTML tag, the following PacketShaper-specific tags can be used on a customer portal page. Notice that all of the PacketShaper tags begin and end with percent signs. If you want to use a percent sign elsewhere on the page, you must use two of them (%%), otherwise the PacketShaper will try to interpret the percent sign as a PacketShaper tag.

Customer Portal Tag	Description
%IF(*<condition>*)% ... %ENDIF%	Checks a condition to see if it's true. If the condition is true, one or more lines of code will be executed; if the condition is false, different lines of code will be executed. Example: `%IF (%PCLASSLIST.ERRORSTATUS% != "")%` `<h3>Errors encountered while preparing this report: ` `%PCLASSLIST.ERRORSTATUS%  ` `The rest of your report will be rendered as well as is possible, but there may be missing sections.` `Please contact Technical Support with this problem.</h3>` `%ENDIF%`
<IMG SRC=/customer/Cgi1?OP.MEAS.GETGRAPH=	Inserts a graph generated by the PacketShaper For details, see "Include Graphs on a Portal Page" on page 1094.

%LOOP%
...
%ENDLOOP%
	Repeats one or more lines of code
	
	Example:
	
	`%LOOP(PCLASSLIST="Inbound;Inbound/HTTP;Outbound")%`
	`%PCLASSLIST.CURRENT[NAME]%`
	`%PCLASSLIST.CURRENT[BYTES.LAST-MINUTE]%`
	`%ENDLOOP%`

%PCLASSLIST.CURRENT[*<me-var>*.*<date-function>*]%	Retrieves data from the current entry; used in a %LOOP% statement where *<me-var>*– one of the measurement engine variables available to the customer portal *<date-function>*– one of the following: LAST-MINUTE, LAST-15-MINUTES, LAST-60-MINUTES, LAST-FULL-HOUR, LAST-24-HOURS, LAST-7-DAYS, TODAY, YESTERDAY, THIS-MONTH, LAST-MONTH Example: `%PCLASSLIST.CURRENT[BYTES.LAST-MINUTE]%`
%PCLASSLIST.CURRENT[NAME]%	Display current class name; used in a %LOOP% statement
%PCLASSLIST.CUSTOMER.MOTD%	Inserts the message-of-the-day text
%PCLASSLIST.CUSTOMER.NAME%	Inserts the name of the logged-in portal user

Customer Portal Tag	Description
%PCLASSLIST.ERRORSTATUS%	Displays the measurement engine's error message when it's unable to retrieve data; can be used inside an %IF% statement to display the message only if there is an error. Example: `<h3>Errors encountered while preparing this report:  %PCLASSLIST.ERRORSTATUS%</h3>`
%PCLASSLIST.NAMED[*<named-class-number>*.*<me-var>*.*<date.function>*]%	Retrieves data from a named class where *<named-class-number>*– the number of the class in your PCLASSLIST. Counting begins at 0; thus, the first class in your PCLASSLIST has a *<named-class-number>* of 0. *<me-var>*– one of the measurement engine variables available to the customer portal *<date-function>*– one of the following: LAST-MINUTE, LAST-15-MINUTES, LAST-60-MINUTES, LAST-FULL-HOUR, LAST-24-HOURS, LAST-7-DAYS, TODAY, YESTERDAY, THIS-MONTH, LAST-MONTH Example: `%PCLASSLIST.NAMED[1.TCP-EFFICIENCY-PCT.LAST-MONTH]%`
%PCLASSLIST.TIMESTAMP.START%	Inserts the date and time of the beginning of the period for which measurement data was tracked
%PCLASSLIST.TIMESTAMP.END%	Inserts the date and time of the end of the period for which measurement data was tracked
%PUSH(PCLASSLIST'"*<class-names>*")%	Defines the traffic classes that are to be reported on the portal page. It must appear at the beginning of the portal page. You can list up to 20 traffic classes; each class name must be separated by a semi-colon, and quote marks must surround the list of class names. Example: `%PUSH (PCLASSLIST="/Inbound;/Outbound;/Inbound/HTTP;/Outbound/HTTP")%`
%PUSH (PCLASSLIST="SEARCH:NAME=*<string>*")%	Dynamically builds a class list based on a *<string>* in a class' name Example: The following statement builds a class list of all the classes that have "HTTP" in the class name: `%PUSH(PCLASSLIST="SEARCH:NAME=HTTP")%` For example: /Inbound/HTTP, /Outbound/HTTP.

Customer Portal Tag	Description
%PUSH (PCLASSLIST="SEARCH:NOTES=*<string>*")%	Dynamically builds a class list based on a *<string>* in a class' comment (notes) field Example: For each class you want to track on the customer portal page, enter the string "PortalPage" in the class' Comment field. Then use the following statement in the customer portal page: `%PUSH(PCLASSLIST="SEARCH:NOTES=PortalPage")%` If you entered "PortalPage" in the Comment field of the /Inbound, /Outbound, /Inbound/HTTP, /Outbound/HTTP, /Inbound/FTP, and /Outbound/FTP classes, the class list would consist of these six classes. If you want to include additional classes later, you don't need to modify the portal page; just add "PortalPage" to the comment field of these classes.
%PUSH (PCLASSLIST="SEARCH:OWNER=*<string>*")%	Dynamically builds a class list based on a *<string>* in a class' owner field Example: A single PacketShaper tracks data for multiple customers. Each customer has a unique set of classes to measure bandwidth usage. In the class' Owner field, enter the customer's login name. Then use the following statement in the customer portal page: `%PUSH(PCLASSLIST="SEARCH:OWNER=%CUSTOMER.NAME%")%` %CUSTOMER.NAME% evaluates to the name of the logged-in portal user. Thus, when CustomerA logs into the customer portal, he will see statistics for all the classes that have CustomerA in the Owner field. When CustomerB log in, she will see statistics for the classes that have CustomerB in the Owner field.

Measurement Variables Available to Portal Pages

Portal pages can include measurement variables that report on throughput, response time, availability, and performance. The tables below list the PacketShaper measurement engine variables and their corresponding portal variable names. Most of the portal variables are identical to those used in the measurement engine. The only difference is that you cannot use the % sign in portal variables make sure you use **-PCT** instead of the % sign. For example, instead of *tcp-efficiency%*, use *tcp-efficiency-pct*.

Measurement Engine Variable	Portal Page Variable Name
app-availability%	APP-AVAILABILITY-PCT
avg-bps	AVG-BPS
bytes	BYTES
tcp-retx-bytes	TCP-RETX-BYTES
license-overflows	LICENSE-OVERFLOWS
network-delay-avg	NETWORK-DELAY-AVG
network-delay-msec	NETWORK-DELAY-MSEC
peak-bps	PEAK-BPS
pkt-exchange-time-avg	PKT-EXCHANGE-TIME-AVG
pkt-exchange-time-msecs	PKT-EXCHANGE-TIME-MSECS
pkt-exchange-time-samples	PKT-EXCHANGE-TIME-SAMPLES
round-trip-time-msecs	ROUND-TRIP-TIME-MSECS
sample-interval-msecs	SAMPLE-INTERVAL-MSECS
server-delay-avg	SERVER-DELAY-AVG
server-delay-msec	SERVER-DELAY-MSEC
service-level%	SERVICE-LEVEL-PCT
service-level-errors	SERVICE-LEVEL-ERRORS
service-level-threshold%	SERVICE-LEVEL-THRESHOLD-PCT
slow-transactions	SLOW-TRANSACTIONS
TCP-efficiency%	TCP-EFFICIENCY-PCT
total-delay-avg	TOTAL-DELAY-AVG
total-delay-msec	TOTAL-DELAY-MSEC
total-delay-threshold	TOTAL-DELAY-THRESHOLD
total-trans	TOTAL-TRANS

Variables used in portal pages do not need to be in upper case.

Pre-Installation Check List

You will need to supply the information described below to configure your PacketShaper. To prepare for Guided Setup, fill in the third column with your settings. You may need to consult your system administrator or network administrator.

Parameter	Description	Your Setting
IP Address	IP address to be assigned to the PacketShaper, for example *10.1.2.3* Consult your network administrator to obtain an unassigned address.	
Net Mask	Subnet mask for the subnet on which the PacketShaper resides, for example *255.255.255.0*	
Gateway	IP address the PacketShaper uses to reach other networks, for example *10.1.2.254* The gateway routes PacketShaper-initiated transactions to a non-local address.	
DNS List *(optional)*	IP address(es) of the Domain Name Server(s), for example *192.181.5.1* You can enter a maximum of eight addresses, separated by spaces. **Caution:** DNS must be configured and working correctly for the licensing process to work.	
Domain Name *(optional)*	Default domain name, for example *bluecoat.com*	

Parameter	Description	Your Setting
Inbound and Outbound Rates	Inbound and outbound link speeds in bits per second. Use the same link speed as that set in the network router. Data rates may be specified as integer bits per second, followed by a “k” (thousands), “M” (millions), or “G” (billions), or specified symbolically (“T1”, “E1”, “T3”). Examples: *1.536 M* or *T1*. ■ For a full-duplex WAN link, enter the total link speed for the inbound and outbound rates. Because full-duplex has wires that can simultaneously communicate in both inbound and outbound directions, you should enter the same rate for Inbound Rate and Outbound Rate. For example, if you have two T1 lines (3 Mbps), you should enter 3M for Inbound Rate and 3M for Outbound Rate. ■ In rare situations in which the PacketShaper is managing half-duplex links, split the rate between the inbound and outbound links. For example, if you are managing a 10 Mbps half-duplex link, you could configure 5 Mbps for the Inbound Rate and 5 Mbps for the Outbound Rate since data can be transmitted in only one direction at a time. **Note:** If you set any NIC in the PacketShaper to half-duplex, you will not be able to transmit and receive at the same time, limiting the bandwidth to half. For example, if you have set 10 Mbps half-duplex on the NIC, the PacketShaper will be able to pass only 5 Mbps IN and 5 Mbps OUT at one time. Hence, you will not be able to manage a WAN link size greater than 5 Mbps. ■ The linksize key may enforce an upper limit on the configurable link size. **Note:** 10BaseT links rarely reach the 10 Mbps limit. Keep Ethernet’s practical limits in mind when configuring rates. ■ If your unit is using NICs to manage different WAN links and you don’t want to control each NIC separately, the rate should be the size of the smallest NIC. For example, if you have two 100 Mbps NICs managing two links, you should specify 100M for the rate. On the other hand, if you want to control each link separately, the rate should be the sum of the link speeds on all devices. For example, if the built-in device is controlling a T1 line (1.5 Mbps) and an option card is managing two T1 lines (3.0 Mbps), you should specify 4.5M for the rate. To control traffic across each link separately, you can create a class for each device (for example, *slot1* and *slot2*) and assign partitions that match the link size (1.5M for the slot1 class and 3.0M for the slot2 class). If your unit is using two NICs to manage a single WAN link, specify the WAN link speed for the rate. Although the Info page will give you an error message (such as “Link speed of 155 Mbps exceeds outside NIC speed of 100 Mbps”) in the latter situation, it is still appropriate to specify the actual size of the link for the rate.	
Packet Shaping	When packet shaping is on, PacketShaper enforces policies and partitions to control traffic. For intial setup, it is recommended to turn shaping off so that you see what type of traffic is running on your network before applying controls. Shaping can be enabled later under the **Setup** tab.	

Parameter	Description	Your Setting
Traffic Discovery	When traffic discovery is on, PacketShaper automatically creates traffic classes based on the traffic it detects.	
Time Zone	A drop-down menu that allows you to select the time zone. (The default is local time, depending on time zone configuration.) Once you configure a time zone, the local time automatically changes at the start and end of daylight savings time.	
Date and Time	Enter the date and time in the format yyyymmddhhmmss. For example*20140106143623*	
Class Tree	**Default**–This bare-bones traffic tree includes: /Inbound and /Outbound classes with a Default class for each, and a Localhost class for the inbound and outbound directions. You can build out this tree by turning on traffic discovery or by manually creating classes for the type of traffic you want to track. **Model**–A pre-configured tree that can be used as is, or modified to suit your needs. This tree organizes network traffic into folders of typical genres, such as VoIP, risky websites, business-critical applications and data, and recreational websites and applications. It includes classes based on service groups or URL categories.	

Priority Guidelines

Use the following table as a guideline for selecting an appropriate priority for a traffic class as part of a priority policy or rate policy. Of course, different applications are of varying urgencies in different environments, so tailor these suggestions to match your own requirements.

Priority	Traffic Characteristics
7	Mission-critical, urgent, important, time-sensitive, interactive, transaction-based.
6	Examples might include SAP, Oracle, a sales website.
5	Important, needed, less urgent and time-sensitive.
4	Examples might include collaboration and messaging systems, such as Microsoft Exchange.
3	Standard service, default, not unusually important or unimportant. Examples might include web browsing.
2	Needed, but low-urgency or large file size. Examples might include FTP downloads, email.
1	Marginal traffic with little or no business importance.
0	Examples might include MP3 music downloads, Internet radio, games.

A traffic class inherits another class' policy if it has no policy of its own. When PacketShaper identifies a class that doesn't have an associated policy, it examines the class' successive siblings, then the parent's siblings, then the grandparent's siblings, and so on, until it finds another matching class with a qualifying inheritable policy. If no other matching class is found, the class will inherit a policy from Inbound/Default or Outbound/Default, which are inheritable classes with a default Priority 3 policy.

Caution: Symantec recommends that you do not change the default Priority 3 policy on the Inbound/Default and Outbound/Default classes.

For more information about policy inheritance, see Policy Inheritance Rules.

Protocol Family

A protocol family represents a single network layer protocol and a number of higher-layer protocols and services.

PacketShaper can detect non-IP protocols only if your router is configured to handle these protocols.

The choices for the Protocol Family field include: **show screen...**

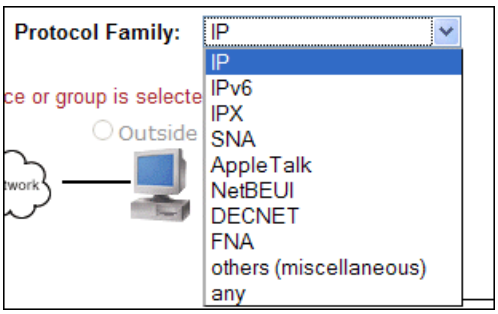

Protocol Family	Classification Match
IP	Services on the IPv4 network layer and TCP and UDP transport layers, for example FTP, HTTP (web), and Telnet.
IPv6	Services using the IPv6 network layer, for example CIFS-IPv6, FTP-IPv6, HTTP-IPv6.
IPX, SNA, AppleTalk, NetBEUI, DECnet, FNA	Non-IP protocol families.
others (miscellaneous)	Services for lesser-known non-IP protocols, such as Spanning Tree and Cisco Discovery. This protocol setting is treated like non-IP protocols, where host, port, service, and location are not applicable.
Any	Matches traffic for any protocol family.

ProxySG Optimization

The ProxySG can perform the following types of optimization:

- **Object caching** – The ProxySG caches HTML pages, images, streaming content, and CIFS file data so that it can serve this data directly to clients; object caching saves time and bandwidth since the content needn't be accessed repeatedly across the WAN.

- **Byte caching** – Replaces byte sequences in traffic flows with reference tokens. The byte sequences and the token are stored in a byte cache on a pair of ProxySG appliances (one at the branch, the other at the data center). When a matching byte sequence is requested or saved, the ProxySG transmits the token instead of the byte sequence. By eliminating repeated patterns of non-cacheable data (or data going over protocols for which a proxy isn't available) from being sent across the WAN, byte caching allows further reduction in WAN bandwidth. The byte cache can be populated by data sent in either direction, and matches can also occur on data flowing in either direction.

- **Compression** – GZIP compression removes extraneous/predictable information from traffic before it is transmitted. The information is decompressed at the destination.

- **Protocol optimization** – Application-layer optimizations use techniques such as read-ahead, pipelining/prefetch, and meta-data caching to reduce "chattiness" in network protocols (such as HTTP and CIFS).

- **TCP optimization** – The ProxySG increases TCP window sizes and improves link throughput across various WAN environments such as MPLS links, satellite links, or congested/lossy links.

RADIUS Accounting Troubleshooting

PacketShaper includes several CLI commands to help you test your RADIUS accounting server.

PING the Accounting Server

To verify that you can connect to the RADIUS accounting server, use the ping command. For example, if the RADIUS accounting server is 10.10.10.10, type **ping 10.10.10.10** at the CLI prompt.

- If you are able to connect to the server, you will see the message "10.10.10.10 is alive."
- If you are unable to connect, you will see the message "no answer from 10.10.10.10."The server may be down or you may have the wrong IP address for the accounting server.

Send Accounting Test Messages

To test and debug the setup of your RADIUS accounting server, use the **radius acct** command in the CLI.

- Type **radius acct start** to send a test message that tells the accounting server that someone logged in.
- Type **radius acct stop** to send a log-off test message.
- Type **radius acct on** to send a message to the server that the RADIUS accounting service is enabled in PacketShaper.
- Type **radius acct off** to send a message to the server that the RADIUS accounting service is disabled in PacketShaper. Note that this command does not affect the setup of the accounting service; if the service was enabled with the **setup radius acct on** command, it will remain enabled (even if you used the **radius acct off** to send a test message that the accounting service is off).

You will get one of the following messages after typing these commands:

Message	What it means	What you should do
RADIUS Accounting request queued	The test message was sent to a queue; the message will be sent to the accounting server from the queue.	Check the RADIUS accounting log to make sure the test message made it from the queue to the server.
RADIUS accounting failed Accounting off	You need to enable the accounting service on PacketShaper.	1. Use the **setup radius acct on** command to enable accounting. 2. Send another accounting test message

Checking the Accounting Log

After sending the accounting test messages, the administrator can then verify that they are in the accounting log. For example, if you sent the four **radius acct** test messages described above, the log would contain entries similar to the following:

```
Thu Aug 2 10:09:16 2001
 User-Name = "radiusAccountingTestUser"
 Acct-Status-Type = Start
 Acct-Session-Id = "radiusTestID"
 NAS-IP-Address = 10.7.51.100
 Acct-Delay-Time = 0
 Client-IP-Address = 10.7.51.100
 Timestamp = 996772156
 Request-Authenticator = None

Thu Aug 2 10:09:27 2001
 User-Name = "radiusAccountingTestUser"
 Acct-Status-Type = Stop
 Acct-Session-Id = "radiusTestID"
 NAS-IP-Address = 10.7.51.100
 Acct-Delay-Time = 0
 Client-IP-Address = 10.7.51.100
 Timestamp = 996772167
 Request-Authenticator = None

Thu Aug 2 10:09:40 2001
 Acct-Status-Type = Accounting-On
 NAS-IP-Address = 10.7.51.100
 Acct-Delay-Time = 0
 Client-IP-Address = 10.7.51.100
 Timestamp = 996772180
 Request-Authenticator = None

Thu Aug 2 10:09:47 2001
 Acct-Status-Type = Accounting-Off
 NAS-IP-Address = 10.7.51.100
 Acct-Delay-Time = 0
 Client-IP-Address = 10.7.51.100
 Timestamp = 996772187
 Request-Authenticator = None
```

If the test entries don't appear in the accounting log, the following situations may be possible.

Situation	What you should do
Server could be down.	Contact your RADIUS administrator to check the status of the RADIUS accounting server. It's a good idea to configure a secondary server to have a backup in case the primary server fails.

Situation	What you should do
Incorrect IP address for the server.	Contact your RADIUS administrator to verify the host name or IP address of the accounting server.
The accounting service may not be enabled on the RADIUS server side.	Contact your RADIUS administrator to verify that the accounting service is enabled on the RADIUS server.
The server may not be configured to work as a PacketShaper client.	For information on configuring the RADIUS server with PacketShaper-specific attributes, see Configure RADIUS Servers.
The LAN may be busy or down.	Check the status of the network.

Checking the Accounting Server Status

You can check the status of the RADIUS accounting server using the **radius show** command. This command tells you whether the server is up or down, and indicates the number of attempts made to connect to the server, the number of successful connections, and the number of connections that timed out.

RADIUS Authentication Troubleshooting

If you are unable to log into PacketShaper after configuring RADIUS authentication, the RADIUS client or server may be incorrectly configured. PacketShaper includes several CLI commands to help you test your RADIUS configuration. (See Access the CLI.)

PING the Authentication Server

To verify that you can connect to the RADIUS authentication server, use the ping command. For example, if the RADIUS authentication server is 10.10.10.10, type **ping 10.10.10.10** at the CLI prompt.

- If you are able to connect to the server, you will see the message "10.10.10.10 is alive."
- If you are unable to connect, you will see the message "no answer from 10.10.10.10.The server may be down or you may have the wrong IP address for the authentication server."

Send a Test Login Request

You can send a test login request using either the **radius login** (if Password Authentication Protocol is used) or **radius chaplogin** (if Challenge-Handshake Authentication Protocol is used) command. You will get one of the following messages after typing the command:

Message	What it means	What you should do
RADIUS Authentication OK	PacketShaper was able to authenticate the user name and password with the RADIUS authentication server.	Try logging in again through the CLI or browser interface.
Authentication turned off	You need to enable the authentication service in PacketShaper.	1. Use the **setup radius auth on** command to enable authentication. 2. Send another test login request
No server configured	The RADIUS authentication service is turned on in PacketShaper, but the server is not configured.	1. Use the **setup radius auth primary** command and specify the Authentication server IP address, port number, and shared secret. 2. Send another test login request.
Access rejected by server	The user name and/or password is invalid.	Contact your RADIUS administrator to verify that you are using the correct user name and password.

Message	What it means	What you should do
Timeout: Unable to obtain a response from server	The RADIUS authentication service is turned on in PacketShaper and the server is configured. This message could be caused by any of the following situations:	
	Server could be down.	Contact your RADIUS administrator to check the status of the RADIUS authentication server. It's a good idea to configure a secondary server to have a backup in case the primary server fails.
	Incorrect IP address for the server.	Contact your RADIUS administrator to verify the host name or IP address of the authentication server.
	The authentication service may not be enabled on the RADIUS server side.	Contact your RADIUS administrator to verify that the authentication service is enabled on the RADIUS server.
	The server may not be configured to work as a PacketShaper client.	For information on configuring the RADIUS server with PacketShaper-specific attributes, see Configure RADIUS Servers.
	The LAN may be busy or down.	Check the status of the network.
Error: Reply didn t contain an access level attribute	The user name and password are valid, but the user wasn't configured with an access level attribute.	Configure the RADIUS server with an access level attribute for this user.

Checking the Authentication Server Status

You can check the status of the RADIUS authentication server using the **radius show** command. This command tells you whether the server is up or down, and indicates the number of attempts made to connect to the server, the number of successful connections, and the number of connections that timed out.

Rate Policy Examples

Using a rate policy, you can ensure that a session gets its fair share of bandwidth. You can also prevent a flow from unfairly impacting other applications. Rate policies are recommended for applications such as web traffic (HTTP), secure web traffic (SSL), email (SMTP and POP3), file transfer (FTP), newsgroups (NNTP), news (EntryPoint), and streaming media (RealAudio, MPEG-Audio, and MPEG-Video).

The following table lists some typical uses of rate policies.

Traffic Class	Recommended Rate Policy	Explanation
Citrix: Inbound Oracle or PeopleSoft over ICA	10 Kbps guaranteed, burstable to link size at priority 7	Citrix applications are mission-critical traffic that need a specific amount of bandwidth to be usable.
Citrix: Inbound email over ICA	0 Kbps guaranteed, burstable to link size at priority 3	Email is not delay sensitive.
FTP	0 Kbps guaranteed, burstable to link size at priority 2	File transfers tend to be bursty, but not time sensitive.
WindowsMedia or NetShow	18 Kbps guaranteed, burstable to link size at priority 4	Streaming media is time-sensitive traffic that need a specific amount of bandwidth to be usable.

Response Time Measurement's Technical Details

PacketShaper tracks the course of a client-server transaction and uses information about a TCP connection to differentiate one portion of the exchange from another.

From its position between the client and server, PacketShaper makes transmission time and size observations during a transaction. Then, PacketShaper incorporates TCP basics to render accurate performance figures.

The following illustration shows the path of a TCP transaction over time. Arrows indicate packets traveling between the client and the server. As you progress down this diagram, you ll see that time increases, with successive events times noted as TN, with T1 representing the first event and T22 representing the last event.

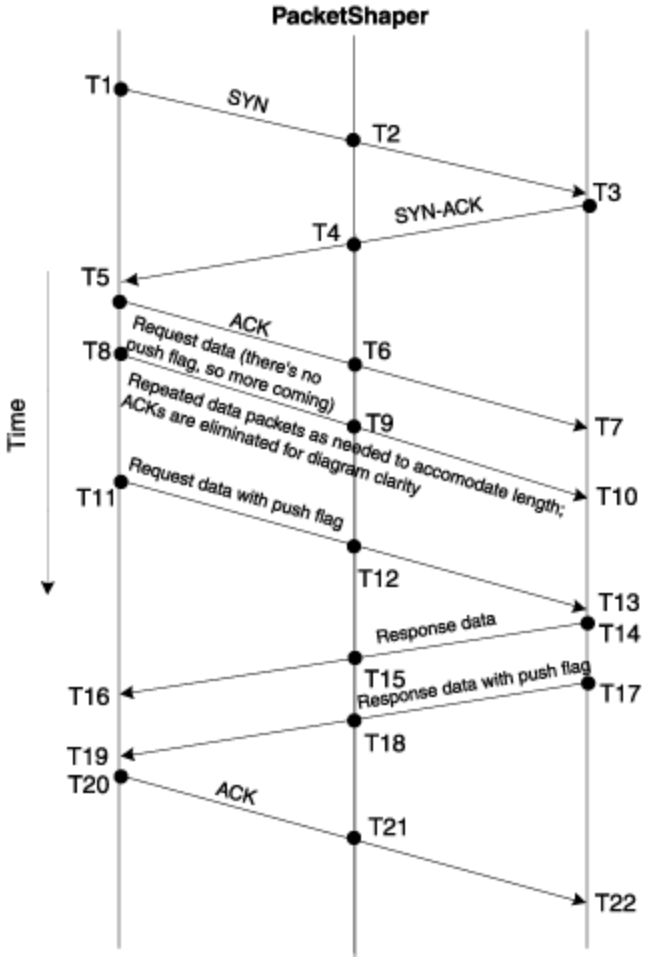

Determining the Round Trip Time (RTT)

A client initiates a server connection with a SYN (synchronize sequence numbers flag) at time T1. PacketShaper notes the SYN at time T2 and passes it along to the server. The server responds with a SYN-ACK (acknowledge flag) at time T3. PacketShaper notes the SYN-ACK at time T4 and passes it back to the client.

TCP stacks usually respond with a SYN-ACK very rapidly, within the kernel and with no context switches. The SYN-ACK follows the SYN almost immediately. Therefore, time T4 minus time T2 gives an accurate measure of the round-trip network delay between PacketShaper and the server. This interchange provides the server transit delay (STD):

```
STD = T4 - T2
```

The client receives the SYN-ACK and issues the final ACK of the three-way handshake at time T5. PacketShaper notes the ACK at time T6, passing it along to the server.

Assume that no processing transpires between the client's receipt of the SYN-ACK and its own corresponding ACK at time T5, as this is handled within the kernel. Time T6 minus time T4 yields an accurate measure of the round-trip network delay between the client and PacketShaper, which provides the client transit delay (CTD):

```
CTD = T6 - T4
```

The server transit delay and the client transit delay yields the total delay between the client and the server for a single round trip:

```
RTT (Round-Trip Time) = STD + CTD
```

Determining the Server Delay

The client initiates its request at time T8, arriving at PacketShaper at time T9. The TCP illustration eliminates the server s corresponding ACKs to simplify the picture, because these ACKs are not relevant to PacketShaper s calculations. The last request packet, sent at time T11, sets its Push Flag to one (1) to indicate that it is the final packet. PacketShaper notes the time of this last request packet at T12.

The last request packet arrives at the server at time T13. The server assembles and processes the request, assembles a response, and sends the first packet (there may be several response packets) at time T14. Time T14 minus time T13 is the actual server-processing time required for the request. These times are not visible to PacketShaper.

PacketShaper knows the server s processing time occurred after PacketShaper saw the last request packet and before PacketShaper saw the first response packet (time T15 minus time T12). PacketShaper also knows that the transit time from PacketShaper to the server and back again (the server transit delay) is another component of this interval.

In addition, a small amount of time is spent serializing the bits in the response packet and preparing them for the bit stream. This time is not included in the original server transit delay because the SYN and ACK packets are extremely small. PacketShaper knows the size of the packet, calculates this preparation time accordingly (D1), and adds it to the server transit delay (STD) before subtracting the sum from the time difference.

Therefore, the server delay is calculated as follows:

```
Server Delay = (T15 - T12) - (STD + D1)
```

Determining the Total Delay

The termination of a transaction is key to calculating the total delay; however, it is not always obvious when a transaction has ended. The combination of the Push Flag from the server and the ACK from the client provides one way of signaling an end, but it is not necessarily the actual end of the transaction. Often applications with long transactions insert Push flags throughout the transaction.

Therefore, in addition to monitoring Push Flags, PacketShaper uses a timer to track transactions. PacketShaper also checks the following conditions to determine if a transaction has truly ended:

- If the server starts sending more data, the timer continues to advance.
- If the client sends a new request, PacketShaper treats this as the end of the last transaction and records the last time noted.
- When there is no activity from either the server or client, PacketShaper considers the transaction complete and records the last time noted.
- When the connection ends, PacketShaper sees the FIN (finish flag) and records the last time noted.

Since the network delay is derived from the total delay, the method for determining total delay impacts the network delay measurements.

In the illustration shown above, PacketShaper notes the last response packet at time T18, and makes sure that it saw all required ACKs for the request packets. It also checks the conditions listed above to determine if the transaction has truly ended.

The client receives the final response packet at time T19 and sends an ACK. The ACK reaches PacketShaper at time T21. The client s perspective of response time starts with sending the first request packet (T8) and ends with receipt of the final response packet (T20). PacketShaper sees that interval as time T9 until time T21. This estimate of the client s view is missing some extra preparation time for serializing the first request packet, assuming that the first request packet is larger than the final ACK. PacketShaper knows the packet-size difference and calculates this small discrepancy (D2).

The total delay is calculated as follows:

```
Total delay = (T21 - T9) + D2
```

Since total delay is a function of transaction size, do not attempt to compare ping times to RTM measurements.

Determining the Network Delay

Using the server delay and the total delay, PacketShaper calculates the amount of time the transaction spent in transit:

```
Network delay = (Total delay) - (Server delay)
```

The network delay reflects all transit time for the transaction. If the transaction s data is large, multiple packets go to and from the server. Only the network delay reflects this overhead. The network delay is not necessarily a multiple of the RTT because multiple packets are not sent consecutively and may overlap to varying degrees.

Troubleshooting Response Time Measurement Issues

At times, you may not be able to set response time measurement settings or view RTM data. This section describes situations that might cause RTM issues.

Accessing Variables After Upgrades

If you upgrade PacketShaper software, you will not be able to immediately access the new response-time variables introduced in the upgraded software version.

To access the variables, reset the measurement engine after installing the upgrade version.

1. Click **Setup** on the navigation bar.

2. From the **Choose Setup Page** list, select **unit resets**.

3. Click **reset measurement data**.

This rebuilds the measurement database to include the expanded set of variables.

> **Note:** Resetting erases all stored measurement data.

Data for Manually Created Classes

If you manually create a traffic class, you must create a corresponding inbound or outbound class to contain the response-time statistics.

Browser Data Differs From Command-Line Data

When you use the browser to view the *Worst Clients* or *Worst Servers* windows, you may see more addresses than when you use the **rtm drilldown** command.

The **rtm drilldown** command is designed to suppress hosts that had fewer than N transactions because they might skew the data. To include every host in the **rtm drilldown** list, set the optional cutoff parameter to 1.

For example:

```
rtm drilldown /outbound/outside/http 10 1
```

Graph Issues

Response-time graphs may not generate. This could be caused by one of the following conditions:

- The measurement engine was recently reset and is in the process of restarting.
- There isn't enough data to graph
- A system error has been identified.

The graphing feature requires a minimum number of data points to generate a line graph, a bar chart, or a pie chart. The number of data points is dependent on the graphing period requested and how recently the measurement engine was reset.

Sample Customer Portal Page with 60 Minute Graphs

This is a sample Customer Portal HTML page with graphs for the preceeding 60 minutes:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE html
PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
"DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en" lang="en">
<head>
<title> Graph sample </title>
</head>
<body>
<h1> Network performance overall, last 60 minutes </h1>
<h2> Inbound throughput </h2>
<IMG
SRC=/customer/Cgi1?OP.MEAS.GETGRAPH=&MEAS.NAME=/Inbound&MEAS.TYPE=link&MEAS.PERIOD=3600&MEA
S.VARLIST=avg-bps%%20peak-bps&MEAS.VARLISTLABEL=AverageRate%%2CPeakRate&MEAS.GRAPHTYPE=line>
<h2> Outbound throughput </h2>
<IMG
SRC=/customer/Cgi2?OP.MEAS.GETGRAPH=&MEAS.NAME=/Outbound&MEAS.TYPE=link&MEAS.PERIOD=3600&ME
AS.VARLIST=avg-bps%%20peak-bps&MEAS.VARLISTLABEL=AverageRate%%2CPeakRate&MEAS.GRAPHTYPE=line>
<h2> Inbound efficiency </h2>
<IMG
SRC=/customer/Cgi2?OP.MEAS.GETGRAPH=&MEAS.NAME=/Inbound&MEAS.TYPE=link&MEAS.PERIOD=3600&MEA
S.VARLIST=tcp-efficiency%%&MEAS.VARLISTLABEL=tcp-efficiency&MEAS.GRAPHTYPE=line>
<h2> Outbound efficiency </h2>
<IMG
SRC=/customer/Cgi2?OP.MEAS.GETGRAPH=&MEAS.NAME=/Outbound&MEAS.TYPE=link&MEAS.PERIOD=3600&ME
AS.VARLIST=tcp-efficiency%%&MEAS.VARLISTLABEL=tcp-efficiency&MEAS.GRAPHTYPE=line>
</body>
</html>
```

Sample Customer Portal Page

This is a sample HTML page for use with the Customer Portal:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE html
 PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
"DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en" lang="en">
<head>
<title> Customer Portal </title>
 %PUSH (PCLASSLIST="/Inbound;/Outbound;/Outbound/HTTP;/Outbound/POP3")%
 %IF (%PCLASSLIST.ERRORSTATUS% != "")%
<h3><b><div align="left"><font color="Red">
 Errors encountered while preparing this report:<br>
 %PCLASSLIST.ERRORSTATUS% <br>
 The rest of your report will be rendered as well as is possible, but there may be missing
sections.
 Please contact Technical Support with this problem.
</font>
</div></b></h3>
 %ENDIF%
</head>
<body>
<h1> <IMG SRC="logo-f.jpg" BORDER="1" ALIGN=TOP WIDTH="640" HEIGHT="75" ALT="Farnsworths
Logo"> </h1>
<h1> <center> Performance report for your subscribed applications </center> </h1>
<h2> <center> Prepared for %PCLASSLIST.CUSTOMER.NAME% </center> </h2>
<h3> <center> For the period ending %PCLASSLIST.TIMESTAMP.END% </center> </h3>
<h4> Message of the day: %PCLASSLIST.CUSTOMER.MOTD% </h4>
<h3> Aggregate Traffic Status </h3>
<p>
 The first table shows the overall performance of your network: throughput in each
direction
 in bits per second,
 and network efficiency, which is the percentage of TCP packets that are not
retransmissions.
</p>
<p align=left>
<TABLE summary="Shows the aggregate traffic status" cellspacing=4>
<TR>
<TH> Metric </TH>
<TH> Status </TH>
<TH> Last minute </TH>
<TH> Last 60 minutes </TH>
<TH> Yesterday </TH>
<TH> This month </TH>
<TH> Last month </TH>
</TR>
<TR>
<TD> Inbound Throughput (bps) </TD>
```

```
<TD align=right BGCOLOR=
 %IF (%PCLASSLIST.NAMED[0.AVG-BPS.LAST-MINUTE]% .GT. "2M")%
"red"
 %ELSEIF (%PCLASSLIST.NAMED[0.AVG-BPS.LAST-MINUTE]% .GT. "1.9M")%
"yellow"
 %ELSE%
"green"
 %ENDIF%
>
 
</TD>
<TD align=right> %PCLASSLIST.NAMED[0.AVG-BPS.LAST-MINUTE]% </TD>
<TD align=right> %PCLASSLIST.NAMED[0.AVG-BPS.LAST-60-MINUTES]% </TD>
<TD align=right> %PCLASSLIST.NAMED[0.AVG-BPS.YESTERDAY]% </TD>
<TD align=right> %PCLASSLIST.NAMED[0.AVG-BPS.THIS-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[0.AVG-BPS.LAST-MONTH]% </TD>
</TR>
<TR>
<TD> Outbound Throughput (bps) </TD>
<TD align=right BGCOLOR=
 %IF (%PCLASSLIST.NAMED[1.AVG-BPS.LAST-MINUTE]% .GT. "2M")%
"red"
 %ELSEIF (%PCLASSLIST.NAMED[1.AVG-BPS.LAST-MINUTE]% .GT. "1.9M")%
"yellow"
 %ELSE%
"green"
 %ENDIF%
>
 
</TD>
<TD align=right> %PCLASSLIST.NAMED[1.AVG-BPS.LAST-MINUTE]% </TD>
<TD align=right> %PCLASSLIST.NAMED[1.AVG-BPS.LAST-60-MINUTES]% </TD>
<TD align=right> %PCLASSLIST.NAMED[1.AVG-BPS.YESTERDAY]% </TD>
<TD align=right> %PCLASSLIST.NAMED[1.AVG-BPS.THIS-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[1.AVG-BPS.LAST-MONTH]% </TD>
</TR>
<TR>
<TD> Inbound Efficiency </TD>
<TD align=right BGCOLOR=
 %IF (%PCLASSLIST.NAMED[0.TCP-EFFICIENCY-PCT.LAST-MINUTE]% .LT. 95)%
"red"
 %ELSEIF (%PCLASSLIST.NAMED[0.TCP-EFFICIENCY-PCT.LAST-MINUTE]% .LT. 98)%
"yellow"
 %ELSE%
"green"
 %ENDIF%
>
 
</TD>
<TD align=right> %PCLASSLIST.NAMED[0.TCP-EFFICIENCY-PCT.LAST-MINUTE]% </TD>
<TD align=right> %PCLASSLIST.NAMED[0.TCP-EFFICIENCY-PCT.LAST-60-MINUTES]% </TD>
<TD align=right> %PCLASSLIST.NAMED[0.TCP-EFFICIENCY-PCT.YESTERDAY]% </TD>
```

```
<TD align=right> %PCLASSLIST.NAMED[0.TCP-EFFICIENCY-PCT.THIS-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[0.TCP-EFFICIENCY-PCT.LAST-MONTH]% </TD>
</TR>
<TR>
<TD> Outbound Efficiency </TD>
<TD align=right BGCOLOR=
 %IF (%PCLASSLIST.NAMED[1.TCP-EFFICIENCY-PCT.LAST-MINUTE]% .LT. 95)%
"red"
 %ELSEIF (%PCLASSLIST.NAMED[1.TCP-EFFICIENCY-PCT.LAST-MINUTE]% .LT. 98)%
"yellow"
 %ELSE%
"green"
 %ENDIF%
>
 
</TD>
<TD align=right> %PCLASSLIST.NAMED[1.TCP-EFFICIENCY-PCT.LAST-MINUTE]% </TD>
<TD align=right> %PCLASSLIST.NAMED[1.TCP-EFFICIENCY-PCT.LAST-60-MINUTES]% </TD>
<TD align=right> %PCLASSLIST.NAMED[1.TCP-EFFICIENCY-PCT.YESTERDAY]% </TD>
<TD align=right> %PCLASSLIST.NAMED[1.TCP-EFFICIENCY-PCT.THIS-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[1.TCP-EFFICIENCY-PCT.LAST-MONTH]% </TD>
</TR>
</TR>
</TABLE>
</p>
<p>
<a href="#Aggregate_Table_Colors"> What the status colors mean </a>
</p>
<p align=left>
<h3> Aggregate Response Time Performance for HTTP </h3>
 This table displays response-time measurements for the traffic class "/Outbound/HTTP",
which represents your
 HTTP traffic.
</p>
<p align=left>
<TABLE cellspacing=4 summary="Shows the Response Time Performance">
<TR>
<TH>   </TH>
<TH> Application <BR> Server (ms) </TH>
<TH> Transport <BR> Provider <BR> Network (ms) </TH>
<TH> Total <BR> Service <BR> Delay (ms) </TH>
</TR>
<TR>
<TD> Last minute </TD>
<TD align=right> %PCLASSLIST.NAMED[2.SERVER-DELAY-AVG.LAST-MINUTE]% </TD>
<TD align=right> %PCLASSLIST.NAMED[2.PKT-EXCHANGE-TIME-AVG.LAST-MINUTE]% </TD>
<TD align=right> %PCLASSLIST.NAMED[2.TOTAL-DELAY-AVG.LAST-MINUTE]% </TD>
</TR>
<TR>
<TD> Last 60 minutes </TD>
<TD align=right> %PCLASSLIST.NAMED[2.SERVER-DELAY-AVG.LAST-60-MINUTES]% </TD>
<TD align=right> %PCLASSLIST.NAMED[2.PKT-EXCHANGE-TIME-AVG.LAST-60-MINUTES]% </TD>
```

```
<TD align=right> %PCLASSLIST.NAMED[2.TOTAL-DELAY-AVG.LAST-60-MINUTES]% </TD>
</TR>
<TR>
<TD> Yesterday </TD>
<TD align=right> %PCLASSLIST.NAMED[2.SERVER-DELAY-AVG.YESTERDAY]% </TD>
<TD align=right> %PCLASSLIST.NAMED[2.PKT-EXCHANGE-TIME-AVG.YESTERDAY]% </TD>
<TD align=right> %PCLASSLIST.NAMED[2.TOTAL-DELAY-AVG.YESTERDAY]% </TD>
</TR>
<TR>
<TD> This month </TD>
<TD align=right> %PCLASSLIST.NAMED[2.SERVER-DELAY-AVG.THIS-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[2.PKT-EXCHANGE-TIME-AVG.THIS-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[2.TOTAL-DELAY-AVG.THIS-MONTH]% </TD>
</TR>
<TR>
<TD> Last month </TD>
<TD align=right> %PCLASSLIST.NAMED[2.SERVER-DELAY-AVG.LAST-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[2.PKT-EXCHANGE-TIME-AVG.LAST-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[2.TOTAL-DELAY-AVG.LAST-MONTH]% </TD>
</TR>
</TABLE>
</p>
<p align=left>
<h3> Aggregate Response Time Performance for POP3 </h3>
 This table displays response-time measurements for the traffic class "/Outbound/POP3",
which represents
 your email traffic.
</p>
<p align=left>
<TABLE cellspacing=4 summary="Shows the Response Time Performance">
<TR>
<TH>   </TH>
<TH> Application <BR> Server (ms) </TH>
<TH> Transport <BR> Provider <BR> Network (ms) </TH>
<TH> Total <BR> Service <BR> Delay (ms) </TH>
</TR>
<TR>
<TD> Last minute </TD>
<TD align=right> %PCLASSLIST.NAMED[3.SERVER-DELAY-AVG.LAST-MINUTE]% </TD>
<TD align=right> %PCLASSLIST.NAMED[3.PKT-EXCHANGE-TIME-AVG.LAST-MINUTE]% </TD>
<TD align=right> %PCLASSLIST.NAMED[3.TOTAL-DELAY-AVG.LAST-MINUTE]% </TD>
</TR>
<TR>
<TD> Last 60 minutes </TD>
<TD align=right > %PCLASSLIST.NAMED[3.SERVER-DELAY-AVG.LAST-60-MINUTES]% </TD>
<TD align=right > %PCLASSLIST.NAMED[3.PKT-EXCHANGE-TIME-AVG.LAST-60-MINUTES]% </TD>
<TD align=right > %PCLASSLIST.NAMED[3.TOTAL-DELAY-AVG.LAST-60-MINUTES]% </TD>
</TR>
<TR>
<TD> Yesterday </TD>
<TD align=right > %PCLASSLIST.NAMED[3.SERVER-DELAY-AVG.YESTERDAY]% </TD>
<TD align=right > %PCLASSLIST.NAMED[3.PKT-EXCHANGE-TIME-AVG.YESTERDAY]% </TD>
```

```
<TD align=right > %PCLASSLIST.NAMED[3.TOTAL-DELAY-AVG.YESTERDAY]% </TD>
</TR>
<TR>
<TD> This month </TD>
<TD align=right> %PCLASSLIST.NAMED[3.SERVER-DELAY-AVG.THIS-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[3.PKT-EXCHANGE-TIME-AVG.THIS-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[3.TOTAL-DELAY-AVG.THIS-MONTH]% </TD>
</TR>
<TR>
<TD> Last month </TD>
<TD align=right> %PCLASSLIST.NAMED[3.SERVER-DELAY-AVG.LAST-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[3.PKT-EXCHANGE-TIME-AVG.LAST-MONTH]% </TD>
<TD align=right> %PCLASSLIST.NAMED[3.TOTAL-DELAY-AVG.LAST-MONTH]% </TD>
</TR>
</TABLE>
</p>
<h2> <center> Appendix </center> </h2>
<a name="Aggregate_Table_Colors">
<p>
 Key to the status colors for the Aggregate Traffic Status table:
</p>
<p align=left>
<table cellspacing=4>
<TR>
<TH> Metric </TH>
<TH> Color </TH>
<TH> Meaning </TH>
</TR>
<TR>
<TD> Throughput </TD>
<TD BGCOLOR="red">   </TD>
<TD align=right> Over 2 million bps </TD>
</TR>
<TR>
<TD>   </TD>
<TD BGCOLOR="yellow">   </TD>
<TD align=right> 1.9 million - 2 million bps </TD>
</TR>
<TR>
<TD>   </TD>
<TD BGCOLOR="green">   </TD>
<TD align=right> Less than 1.9 million </TD>
</TR>
<TR>
<TD> Network efficiency </TD>
<TD BGCOLOR="red">   </TD>
<TD align=right> Less than 95 percent </TD>
</TR>
<TR>
<TD>   </TD>
<TD BGCOLOR="yellow">   </TD>
<TD align=right> 95 - 98 percent </TD>
```

```
</TR>
<TR>
<TD>   </TD>
<TD BGCOLOR="green">   </TD>
<TD align=right> Over 98 percent </TD>
</TR>
</table>
</a>
</p>
</body>
</html>
```

Secure Logins

The browser interface uses HTTPS (HTTP over Secure Sockets Layer) protocol and the CLI uses SSH (Secure Shell) protocol.

Browser Interface

By default, HTTP access to PacketShaper is disabled so you must log in with HTTPS. This will ensure a secure login to the browser interface.

- Type the URL **https://*<ip address>***. When the Login screen appears, the **Secure Login** check box will be selected automatically.

If you want the ability to log in to the browser without a secure connection, you can enable HTTP access. See " Specify Security Settings " on page 297.

Logging in Securely the First Time

The first time you access PacketShaper with a secure connection, a security alert will appear. This message appears because the security certificate was generated by Symantec, not a Certificate Authority (a trusted third-party organization that issues digital certificates) that is already preconfigured in your browser.

Command-Line Interface

For a secure login to the CLI, use any SSH client, such as SecureCRT for Windows or OpenSSH for UNIX operating systems.

Windows SecureCRT Settings

If you are using a Windows SecureCRT SSH client to connect with the PacketShaper, make sure the settings are selected as follows.

Protocol: SSH v2
Hostname: : IP address of PacketShaper
Port: 22 (unless changed by user)
Username: can be left blank
Cipher: 3DES
Authentication: Password
SSH Server: Standard (the default setting is **DataFelllows 2.0.13** this setting will not work and must be changed to **Standard**)

Security Alert

If you are using PacketShaper software in Internet Explorer on a secure connection (HTTPS), you may occasionally get the following security alert:

```
This page contains both secure and non-secure items. Do you want to download the non-secure items?
```

Internet Explorer generates this message if a secure Web page uses frames and one of the frames doesn't load, either because you canceled the load or because the load request didn't work. (For additional information, see the article on the Microsoft support site.)

If you get this message in PacketShaper, here is what you can expect:

- If you answer **Yes** to download the non-secure items, the page will display but the Secure Lock indicator will disappear from the status line.
- If you answer **No**, you will see the message "Navigation Canceled" and the Secure Lock indicator will remain in the window's status line.

Regardless of the way that you respond to the security alert, you can continue using PacketShaper and data will still be transmitted securely.

Selecting a Time Period

The ability to create reports for a specific time period is useful for creating comparison reports. For example, you can print one report that analyzes traffic over the last week (a period of one week, with **now** as the end date and time), and then print a second report that analyzes traffic for the week before that (a period of one week, with a specific date specified for the end date the last day you want to analyze).

If you want to view a report of the last hour (the default), eight hours, day, or week, just use one of the preset buttons:

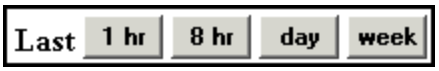

Or, to specify a different time period:

1. In the **Show** field, specify the number of minutes, hours, days, weeks, or months for which you want traffic analyzed.

For example, by entering **2** in the **Show** field and selecting **week** for the unit of time, you create a report with the data PacketShaper collected over the last two weeks. The minimum period you can specify is 60 minutes (or 1 hour); the maximum is 60 days (or 2 months).

2. To choose a specific end date for the report, select the month and day from the **Before** drop-down lists. If you want the current date as the end date, leave the default values of **now**.

> **Note:** Although you can select a date from the previous 12 months, the unit may not store data that far back the amount of historical data that is stored depends on how many traffic classes you have in your tree. The unit stores measurement data for a minimum of 31 days in hourly intervals and at least 1 day of data in one-minute intervals.

3. To choose a specific end time for the report, select the time from the drop-down list. If you want the current time as the end time, leave the default values of **now**.

4. Click **update** to update the statistics for the new time period.

> **Note:** To update the summary at regular intervals, select an interval from the drop-down list located next to the **update** button.

5. To print the report, click **print**.

Service Groups Best Practices

Be sure to follow these best practices when using the service group feature.

Initial Setup of Service Groups

- Look at the default set of service groups and determine whether these groupings serve your monitoring and/or control requirements. If necessary, create your own custom groups and move services between groups as required.

- It's important to note that in order to classify by service group, you need to manually create classes based on each service group. Service group classes are not created automatically via traffic discovery.

- If you already have a service-based traffic tree that was built through auto-discovery and you want to switch to a traffic tree that is based on service group classes, you should:

    1. Turn off global traffic discovery.

    2. Reset the class tree.

        Note: Measurement data associated with the deleted classes will be lost.

    3. Create classes for each service group you want to monitor and/or control, and turn on class discovery if you want child classes created for each service that is discovered in the service group.

    4. Turn on global traffic discovery.

- New installations should follow the same steps described above, except step 2 is unnecessary.

Deleting Custom Groups

- Before deleting a custom group, delete any classes you may have created based on this service group. If you don't delete the class first, it will have a configuration error because the service group on which it is based no longer exists.

- When you delete a custom group, all services are moved into the Unassigned group. You can then return these services back to their original group(s), if you like. Services may also be moved into the Unassigned group when you re-inherit a local overridden service group from the parent configuration.

Service Group Maximum

The maximum number of service groups allowed on a PacketShaper is 50. There are 25 built-in groups plus you can create up to 25 custom groups.

Browser Refresh may be Required

Occasionally the configuration may not be updated to reflect service group changes. If this happens, refresh your browser. For example, a browser refresh may also be required when service groups are re-inherited with the **re-inherit all** button.

SNMP Trap Filtering

When SNMP is enabled and configured on the PacketShaper, the appliance will send SNMP traps for a variety of component failures or status changes. See "PacketShaper SNMP Traps" on page 1088 for a complete list. However, if you aren't interested in being notified about a particular situation, you can create a *trap filter*.

To prevent specific traps from being sent:

1. Determine the OIDs of the traps you aren't interested in monitoring. You can do a web search of the trap name to find the OID. For example, search for *linkDown trap OID* to find its OID (1.3.6.1.6.3.1.1.5.3).
2. Create an SNMP view that excludes the OID you aren't interested in monitoring. For example, create a view called TrapFilter, add the OID 1.3.6.1.6.3.1.1.5.3, and select exclude. See "Create a New View" on page 337.
3. To the same view created in step 2, add any other OIDs you want to exclude from notification. See "Add Additional OIDs to a View" on page 338.
4. Associate this view with the SNMP notify target. See "Create a New Notify Target " on page 349.

Supported Ciphers and Algorithms for SSH and SSL

SSH

PacketShaper supports SSHv2 protocol; SSHv1 is less secure and not supported. However, not all ciphers and MAC (Message Authentication Code) algorithms are supported for these protocol versions. A client can select any one of the supported ciphers listed below that are also supported by the remote login utility.

Supported Ciphers for SSHv2

aes128-ctr

aes192-ctr

aes256-ctr

aes128-cbc

aes192-cbc

aes256-cbc

Supported MAC Algorithms for SSHv2

hmac-sha2-256

hmac-sha2-512

Unsupported Ciphers

The following ciphers have been disabled due to weakness:

All SSHv1 ciphers (for example: des, des-dbc, blowfish, idea)

aes128-cbc

aes192-cbc

aes256-cbc

aes128-ctr

aes192-ctr

aes256-ctr

arcfour

arcfour128

arcfour256

blowfish-cbc

cast128-cbc

Unsupported MAC Algorithms

The following MAC algorithms have been disabled due to weakness:

hmac-md5

hmac-md5-96

hmac-ripemd160

hmac-sha1

hmac-sha1-96

SSL

Supported Ciphers for SSL

PacketShaper 11.5.3 and higher versions include the secure cipher suite for Transport Layer Security (TLS) 1.2. The following ciphers are supported:

TLS_DHE_RSA_WITH_AES_128_CBC_SHA256

TLS_DHE_RSA_WITH_AES_128_GCM_SHA256

TLS_DHE_RSA_WITH_AES_256_CBC_SHA256

TLS_DHE_RSA_WITH_AES_256_GCM_SHA384

TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256

TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256

TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384

TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384

TLS_RSA_WITH_AES_128_CBC_SHA256

TLS_RSA_WITH_AES_128_GCM_SHA256

TLS_RSA_WITH_AES_256_CBC_SHA256

TLS_RSA_WITH_AES_256_GCM_SHA384

PacketShaper Syslog Warn Messages

At the "warn" level, PacketShaper will send the following messages to the syslog server:

login failed for *<userSessionUserName>*

login failed

Touch access failed

Hard drive is down

Hard drive is back up

Error starting measurement engine -- not started

Characterizing Traffic

For each type of traffic you want to manage, consider its behavior with respect to four characteristics: importance, time sensitivity, size, and jitter.

Each characteristic below has an associated question to ask yourself, as well as several examples of applications or protocols that fit the definition for *a YES* answer and that fit the definition for a *NO* answer.

Importance: ***Is the traffic critical to organizational success?***

Yes, Important

- SAP to a manufacturing business
- Quake to a provider of gaming services
- PeopleSoft to a support organization
- Email to a business

No, Not Important

- Real Audio to a non-related business
- Games in a business context
- Instant messaging in a classroom

Time Sensitivity: ***Is the traffic interactive or particularly latency sensitive?***

Yes, Urgent

- Telnet
- Interactive applications running over Citrix
- Oracle

No, Not Urgent

- Print
- Email
- File transfers

Size: ***Are flows large and bandwidth hungry, expanding to consume all available bandwidth?***

Yes, Large and Demanding

- Music downloads
- Email with large attachments
- Web browsing

No, Small

- Telnet
- ICMP
- TN3270

Jitter: ***Does the traffic require smooth consistent delivery or it loses value, suffering stutter and static?***

Yes, Sensitive to Jitter

- VoIP
- WindowsMedia
- Real Audio
- Distance-learning applications

No, Oblivious to Jitter

- Email
- Print
- MS SQL
- TN3270

These characteristics are useful when determining what types of policies and partitions are appropriate. For example, if an application's traffic is large and demanding, a capped partition is probably in order, independent of whether the application is important or not.

See Control Strategy Overview for more information about approaching control features for the first time. See Policy and Partition Guidelines to use these four characteristics in choosing an appropriate policy or partition for an application.

Traffic Tree Designs

The way PacketShaper organizes traffic dictates how it can analyze and control that traffic. Therefore, your traffic tree configuration is a crucial choice in determining the features you have available.

For example, if your tree doesn't distinguish SAP from Oracle, then you can't measure distinct response times for each. Or, if your tree doesn't distinguish traffic to your Paris office from traffic to your Oslo office, then you won't be able to compare traffic volumes or give prescribed amounts of bandwidth to each branch.

Traffic trees are very flexible. You can customize your own trees' classes using many types of criteria including application, location, protocol, URL categories, host lists, addresses, bit markings such as ToS, and many other strategies. We suggest a few options of tree types with their implications here. You are not limited to these trees, but you can use them in initial strategies.

Application-Based Tree
Simple Location-Based Tree
Location-Based Tree with Per-Location Applications
Location-Based Tree with Global Applications

To help in determining which traffic trees might suit your own circumstances, have the answers to these questions ready before reading the tree descriptions.

- How many branches do you manage?
- How many applications (or other subsets of traffic, such as all traffic to one particular Oracle database) do you want to manage separately? For each remote office?
- Can you fit your applications into 3 to 5 categories of management strategies? Would it be okay if you had usage statistics and graphs only for each category (Critical, for example) instead of for each application (SAP, for example)?
- Do you want separate usage statistics and graphs for each application? For each application at each branch?
- Do you want to control each application for each branch distinctly? Or can one policy/partition strategy apply to each application, no matter the application's branch? For example, does Citrix need a different policy at one branch than at another? Perhaps because the different branches' link capacities vary?

> **Note:** You can find instructions for how to create multiple types of traffic trees in Classify Network Traffic.

Application-Based Tree

show tree example

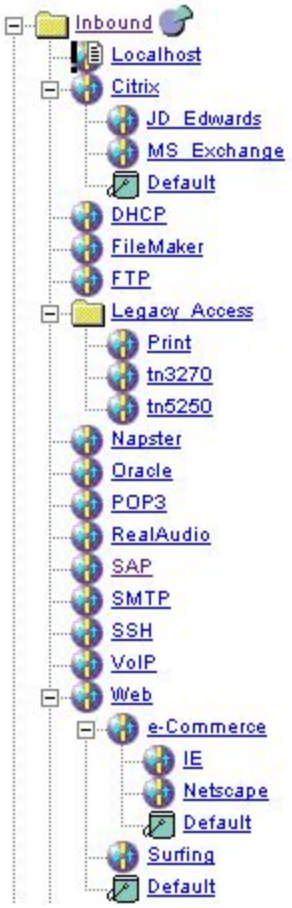

This tree is for a WAN or Internet link at any main or branch site where you want to vary traffic management strategies according to type of traffic (such as application) rather than according to destination or location.

An application-based tree is most common at branch sites and main sites' Internet links. Automatic traffic discovery and classification are a good basis for the traffic classes, but you can also add or tailor classes based on other criteria. It's a simple tree to create and use.

Two variations of the application-based tree are:

- Category-based tree: monitors and controls traffic based on URL categories.
- Model tree: a pre-configured tree that organizes similar traffic into folders.

Capabilities

- The PacketShaper offers all its monitoring and analysis features for any traffic class.
  For example, you can compare bandwidth usage for SAP and Oracle, but not for Sao Paulo and Bogota.
- The PacketShaper with the shaping module offers all its control features.
  For example, you can allocate 512 Kbps to all of SAP and a steady 22 Kbps to each VoIP session.

Limitations

No insight into or control over per-location behavior.

Scaling Considerations

Choose a product model based on your link's capacity, anticipated maximum number of concurrent flows, and/or number of users. This configuration rarely bumps into limits on traffic classes or matching rules.

There is less growth-associated overhead than for location-based trees. Automatic traffic discovery detects many new applications and creates classes for you.

Simple Location-Based Tree with Locations Only

show tree example

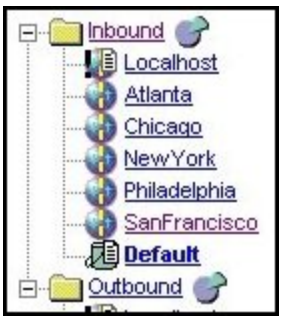

This tree is for a main site's WAN or Internet link with traffic that goes to multiple branches. It categorizes first by travel direction and then by location. It's appropriate for occasions when your primary concern is to provision bandwidth, and you don't particularly care how it's used or how applications perform. Additionally, it can be appropriate for a main site in topologies where other PacketShapers manage applications at each location, and the main site's PacketShaper just manages the amount of traffic going out to each branch.

Capabilities

- The PacketShaper offer all its monitoring and analysis features - but, on a global and per-location basis only.
  For example, you can compare bandwidth usage in Sao Paulo with usage in Bogota, but you cannot see how much of

Bogota's usage is music downloads or SAP.

- The PacketShaper with the control module offers all its control features - but, on a global and per-location basis only. For example, you can allocate 512 Kbps to Bogota, but you cannot block music downloads or protect SAP.

Limitations

No insight into or control over per-application behavior. No visibility into which applications travel your network.

Complementary Products

IntelligenceCenter can collect additional information that allows you to drill down into detailed data for traffic flows. With Flow Detail Records enabled on your PacketShaper, IntelligenceCenter can collect and display data not visible from the PacketShaper alone, including the traffic's service or application, number of retransmitted bytes, and VoIP statistics for VoIP streams implementing the RTCP standard.

Scaling Considerations

This management option scales well, as it accommodates many locations easily. Choose a product model based on your link's capacity, anticipated maximum number of concurrent flows, and number of traffic classes to accommodate your number of locations (times 2, to have each location under both Inbound and Outbound).

There is little overhead involved in adding branches, as you need only to create an additional class under both the Inbound and Outbound classes.

Configuration Suggestions

The most common criteria to define location-based classes are subnets or host lists. You can create initial location classes one-by-one by hand, or, if there are many classes and it gets to be a tedious task, you can automate the process with a command file and a data entry form.

At a main site, give each branch office a static partition with a size equal to the branch's WAN-link capacity. If you are concerned with each user's bandwidth, define a dynamic, per-user partition under each branch's static partition.

Location-Based Tree with Per-Location Applications

show tree example

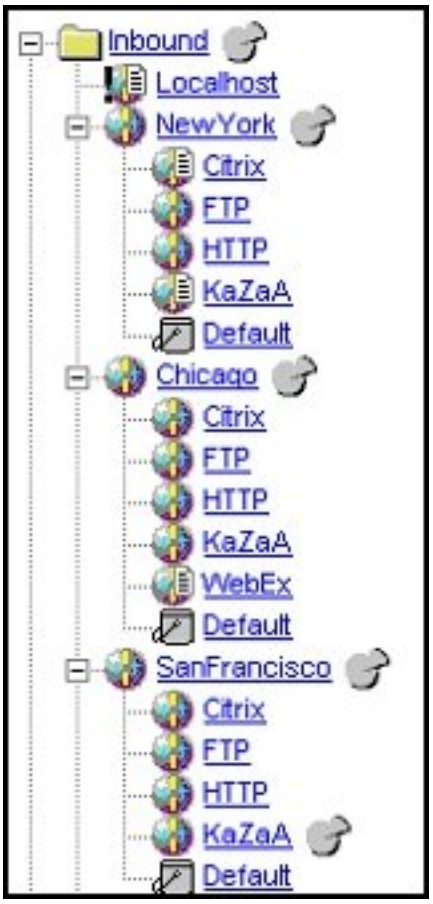

This tree is for a main site's WAN or Internet link with traffic that goes to multiple branches. A location-based tree with per-location applications gives you the most insight and the most control over your applications. But it imposes scaling considerations and takes more time to configure. It categorizes first by travel direction, then by location, and finally by application.

Capabilities

- The PacketShaper offer all its monitoring and analysis features.
  This is an appropriate configuration choice when you want the most detailed level of analysis.
  For example, if you want to know how much New York traffic you have, and you want to know how much SAP traffic you have, and you want to know how much SAP traffic you have in New York then this traffic tree is your best bet.
  In addition, IntelligenceCenter can offer more detailed as well as high-level roll-up reports for combinations of locations.

- The PacketShaper with the control module offers all its control features.
  This is an appropriate configuration choice when you need to control each application to each location distinctly.

For example, suppose the Sao Paulo branch has more capacity than Bogota has. You don't want to limit Sao Paulo's FTP to match the level that Bogota can handle. But you also do not want to burden Bogota with FTP just because Sao Paul can handle it. With this configuration, Sao Paulo and Bogota can each have different FTP limits.

Limitations

You must be prepared to manage the details of each location separately. You must assign appropriate policies for each application you wish to manage at each location.

Scaling Considerations

This style of traffic tree tends to use up many traffic classes: (number of applications+1(for default)) * (number of branches+1) * 2 (for inbound and outbound).
For example, if 15 branches each have 20 applications, 16 * 21 * 2 = 672 classes. You would need a PacketShaper to accommodate this tree, even if your WAN speed doesn't require the capacity.

If you need or will shortly need more traffic classes than are available, you have outgrown a main-site-only topology. Instead, choose a distributed topology with a simple location-based tree type at your main site.

If you add a location, you add just the one class under Inbound and Outbound. Then you can use automatic discovery features (or you can copy other branches' classes).

If you add an application to some or each location, automatic discovery will probably find it. Or, you can copy a manually created class.

You must assign appropriate policies and partitions to each of an application's many classes.

Configuration Suggestions

At a main site, give each branch office a static partition with a size equal to the branch's WAN-link capacity. Then, under each branch, assign partitions and policies to each application that needs protection or that undermines other applications' performance. Use minimum and maximum sizes that fit appropriately under the associated branch's partition. For example, if New York has a 2 Mbps partition, then you could give SAP 500 Kbps, but not 3 Mbps.

Location-Based Tree with Global Applications

show tree example

This tree is for a main site's WAN or Internet link with traffic that goes to multiple branches. A location-based traffic tree with global applications categorizes traffic first by travel direction, then by location. That's all. The metrics, graphs, and reports are the same as those available in the simple location-only tree. The big difference is that you can define separate bandwidth-allocation policies for each application. But all traffic from one application is managed with one strategy, independent of branch.

This option is most common when all branches have the same link size and similar business priorities because each application is managed the same way for all locations. This traffic tree is the most difficult to understand conceptually, but the easiest to configure and needs the least amount of traffic classes, matching rules, and configuration time. It provides a very scalable and easily managed main-site tree for traffic shaping but does limit visibility.

This type of tree is not appropriate for a PacketShaper without the control module, as the only purpose of the application classes at the bottom is to apply policies.

Capabilities

- PacketShaper offers all its monitoring and analysis features but, on a global and per-location basis only.
  For example, you can compare bandwidth usage in Sao Paulo with usage in Bogota, but you cannot see how much of Bogota's usage is Citrix. You cannot see which are the top websites.
- PacketShaper offers all control features on a global and per-location basis, and it offers policies on a per-application basis.
  For example, you can allocate 512 Kbps to Bogota, and you can assign TN3270 a priority of 7. But you cannot assign TN3270 a priority 7 for Bogota and a priority 3 for Quito.

Limitations

No per-application insight. Must use inheritable policies in PacketShaper, one of the product's more conceptually difficult constructs.

Scaling Considerations

This tree scales well, as it accommodates many locations and applications easily.

The number of traffic classes this type of tree requires: ((number of branches) + (number of applications)) * 2 (Inbound and Outbound).
For example, if you have 20 branches and 40 applications, (20 + 40) * 2 = 120 traffic classes, well within the means of all models. You would use link speed or another option to choose your model.

There is little overhead involved in adding branches. You need only to create an additional class under both the Inbound and Outbound classes. Automatic discovery will probably take care of adding applications. And if not, you need only to create a traffic class once for each application under both Inbound and Outbound, as the classes are not repeated under each branch.

Inheritance

A summary of how inheritance applies in this type of tree follows. (See Policy Inheritance Rules for a more complete description of inheritance and examples, if desired.)

PacketShaper classifies each passing traffic flow in the class that matches the flow's branch location. PacketShaper adds the metrics for the traffic flow only to the matching branch class. It applies any partition assigned to the matching branch class. But when PacketShaper finds no policy assigned to the branch class, it continues searching until another matching class with a policy is found.

If the tree is configured correctly, PacketShaper finds an application class as the second match. The flow inherits the second class' policy. But the second class gets no performance metrics recorded other than an additional policy hit.

For example, consider the following tree fragment with three branches and three applications:
Inbound
  London (Partition of 512 Kbps)
  Paris (Partition of 512 Kbps)
Rome
  SAP (Rate Policy 6)

WebBrowsing (Rate Policy 3)
TN3270 (Priority Policy 7)

When a 100 KB TN3270 flow heads to London, the London class gets those bytes added to its usage figures, the flow is paced to fit within London's partition, and the flow gets top access to the partition's capacity with its inherited priority 7 policy. But the TN3270 class' usage is zero because the London class was the initial classification match. The TN3270 class' policy hits does increase by one. You can generate insightful reports about the London class as a whole, but not about the TN3270 class.

Configuration Suggestions

At a main site, give each branch office a static partition with a size equal to the branch's WAN-link capacity. Do not assign a policy, as traffic will inherit its policy from the application classes. Then, create inheritable classes (beneath the list of locations) for each application and assign appropriate policies.

URL Syntax

When creating a class or matching rule based on a URL, specify the URL according to the syntax rules listed below:

- The legal characters for URLs are 0-9, A-Z, a-z, and the following special characters: $-_.,+!* ()
- Specify only the portion of the URL that follows the server hostname (http://hostname). The hostname goes into the Name field (in the Inside or Outside column in the window, depending on where the web server resides). For example, to specify a remote website http://www.mycompany.com/FavoriteStuff, enter *www.mycompany.com* into the Name field of the Outside column and enter */FavoriteStuff* in the URL Criterion field.
- PacketShaper cannot classify URLs with illegal characters (such as a space). However, when a web browser requests resources from a server using URLs, many special characters are converted to an encoded form before being transmitted, and you can use the encoded form in your URL criterion. The encoded form is a percent character (%) followed by the hexadecimal representation of the character. For example, the space would be sent as %20. You must use the encoded form for these characters when entering URL criterion. For example, to specify the URL*/local contacts/dir1/**, enter */local%20contacts/dir1/** in the URL Criterion field.

  To determine the encoded form for a special character, look up the two-character hex code in an ASCII table and precede it with a percent sign. The encoded forms of commonly-used special characters are listed below:

space	%20
:	%3A
\	%5C

- The asterisk wildcard character (*) is supported.

  For example, you can configure a traffic class to classify GIF files by specifying **.gif* as the URL.
- Slashes are *not* treated as special characters.
- To match a specific URL, include the leading slash.
- Only one URL can be specified per matching rule.
- If some browsers, users, or incoming links are escaping characters that others do not escape, create a matching rule for each version of the URL.

> **Note:** PacketShaper examines only the first 128 bytes of the URL in a flow.

Note: Consider that users are able to arbitrarily convert some or all of the characters in a URL to escaped hexadecimal values. For example the following URL will link you to Yahoo.com http://%79%61%68%6f%6f%2e%43%4f%4d In this example the word yahoo was encoded in lower case while COM is in upper case. The domain name portion of a URL is not case sensitive. In order to classify URLs that contain escaped characters you must create a matching rule containing the same characters, or use a wild card.

Use PacketShaper MIBs

The PacketShaper measurement engine collects data on the traffic flowing through the PacketShaper. You can access this information using your SNMP network management application to poll the appliance.

> **Note:** Delete any prior occurrences of PacketShaper MIBs from your MIB browser/compiler before loading these revised MIBs; if you don't, there will be a conflict of MIB modules.

To use PacketShaper MIBs:

1. Go to the Symantec Support site, MySymantec.com, and log in using your MySymantec credentials.
2. In the products list, locate the serial number of the appliance.
3. Download the MIB zip file for the desired version.
4. Extract the MIBs from the zip file:

   BLUECOAT-PACKETSHAPER-S-SERIES-MIB

   BCSI-MIB

   BLUECOAT-PACKETSHAPER-S-SERIES-RTM-MIB
5. Create a repository for all the MIBs in ~/.snmp/mibs. Copy the aforementioned Symantec MIBs into the repository along with SNMPv2-CONF, SNMPv2-SMI, and SNMPv2-TC. Copy additional MIBs in as needed.
6. Use your network management application's MIB integration tool to compile both PacketShaper MIBs.
7. Query the PacketShaper MIB objects, using your network management application.

Examples

To walk the S-Series MIB with SNMPv2:

```
snmpwalk -mALL -v2c -c <read_community_string> <PacketShaper_IP> psCommonMibS
```

To walk the S-Series MIB with SNMPv3:

```
snmpwalk -mALL -v 3 -n "" -u v3User -l AuthPriv -a MD5 -A <password> -x DES -X <password> <PacketShaper_IP> psCommonMibS
```

To view the structure of the S-Series MIB:

```
snmptranslate -M+. -mBLUECOAT-PACKETSHAPER-S-SERIES-MIB -Tp -IR
```

To set up snmptrapd to receive traps from our remote user:

Add the following to /etc/snmp/snmptrapd.conf where the -e option specifies the localSnmpId of the PacketShaper:

```
createUser -e 0x80001F8804343931323330303031 36 r5 MD5 "password" AES "password"

authuser log r5
```

Then run snmptrapd so it logs to stdout:

```
sudo snmptrapd -f -Lo -c /etc/snmp/snmptrapd.con
```

SNMP GET Limitation

When a user performs an SNMP GET operation to request information on an object with index 0 (such as a class or partition with an ID of 0), the client will not generate the request and the following message appears:

```
Unknown Object Identifier (Index out of range: 0)
```

This occurs when using non-community string based authentication because the snmpget utility considers an index value of 0 in a table to be invalid. To work around this limitation, you can use the **-Ir** option to disable the checking of objects against the range specified in the MIB. For example, to request information about linkTotalRxPkts, use a command similar to the following:

```
snmpget -l NoAuthNoPriv -u <user> -Ir -m all <PacketShaper_IP> 1.3.6.1.4.1.14501.11.2.1.1.2.1.37.0
```

The following result appears:

BLUECOAT-PACKETSHAPER-S-SERIES-MIB::linkTotalRxPkts.0 = Counter32: 0

actionfile sync

For units in shared mode only

Issue this command from an individual PacketShaper to immediately download adaptive response action files prescribed for the unit s PolicyCenter configuration. This command is only required when the PolicyCenter configuration prescription mode has been set to **scheduled** with theactionfile subscribe command.

> **Note:** It is not necessary to issue this command if the prescription mode is currently in its default state, or has been set to **asap** with the **actionfile subscribe** command.

`actionfile sync <seconds>`

If you include the optional *<seconds>* value, the **actionfile sync** operation runs for the specified number of seconds.

agent action

Delete an adaptive response action file, temporarily disable or reenable an existing action file, or modify the value of an existing parameter. Note that this command will not create a new action file, or add a new parameter to an existing action file.

`agent action <name> green|red [on <filename>]|[off]|[delete]|[parm <parm-name> <parm-value>]|[resetparms]`

<name>	Name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example My Agent. If the agent name is a single word, the quotation marks are not necessary.
green|red	Action file will trigger when the green or red threshold is crossed
<[on *<filename>*] | [off] | [delete]	Specify one of the following: **on**: Enables the action file. Specify the name of the action file you want to associate with the agent with the *<filename>* variable. **off**: Disables the action file **delete**: Deletes the action file specification for the agent. The action file is no longer associated with the agent, but the action file is not removed from the unit or PolicyCenter server.
[parm *<parm-name><parm-value>*]	Specify the following: *<parm-name>*: The name of the action file parameter being modified *<parm-value>*: The new value of the parameter
[resetparms]	Specify this operation only when action file parameters have been edited and need updating. agents will not recognizes new action file parameters unless the action file is reset with this variable.

Before you can issue any other agent actionfile commands, you must first issue the command
agent actionfile <name> green|red on *<filename>*
to associate an action file with the agent. You may then issue any of the following commands (see the table above for an explanation of variables):

agent action *<name>* green|red off

agent action *<name>* green|red delete

agent action *<name>* green|red resetparms

agent action *<name>* green|red parm *<parm-name>* *<parm-value>*

For Example:

agent action "Packet Drops" green on actnfile.cmd

agent action "Packet Drops" green parm ClassName /outbound/Citrix

Command Change History

Release	Modification
11.5.1	Command introduced

agent createdefaults

Recreate the default set of agents. The adaptive response feature must be enabled before you can create default agents with this command.

```
agent createdefaults
```

Note that this command will not overwrite any existing default agents that you may have customized, nor does it remove any new agents you may have created.

Command Change History

Release	Modification
11.5.1	Command introduced

agent delete

Delete an existing adaptive response agent. Scoring and status information for the agent will no longer appear in the agent pop-up window on the unit's **info** page or the PolicyCenter **configurations** page.

```
agent delete <name>
```

where *<name>* is the name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example, *High Bandwidth Host* . If the agent name is a single word, the quotation marks are not necessary.

Command Change History

Release	Modification
11.5.1	Command introduced

agent interval

Set an evaluation interval for an adaptive response agent, in minutes. An evaluation interval determines how often the agent checks the status of its target.

```
agent interval <name> <interval> | default
```

where *<name>* is the name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example, *High Bandwidth Host* . If the agent name is a single word, the quotation marks are not necessary. Specify the interval in minutes, or enter **default** for the default evaluation interval. The maximum evaluation interval allowed is 99999 minutes; the minimum is 1 minute.

Command Change History

Release	Modification
11.5.1	Command introduced

agent new

Create a new adaptive response agent based on one of the agent templates. Note that this command creates a new agent, but does not allow you to specify parameter values. Once you have created a new agent, issue the command agent parm to change the parameter values from their default settings. Each PacketShaper or PolicyCenter configuration can have a maximum of 32 agents.

> **Note:** Some agent templates do not allow multiple instances. If you want to create a new agent from the following templates, first delete the existing agent from that template from your unit or PolicyCenter configuration.

- High Bandwidth Host
- New Application
- High Bandwidth New App
- Memory Allocation
- Unit Limits
- System Load

```
agent new <name> <template>
```

<name> Name you want to assign to the agent. An agent name can have up to 32 alphanumeric characters, including -, _, and . (period). If the agent name has a space, the words must be entered within quotation marks, for example, My Agent. If the agent name is a single word, the quotation marks are not necessary.

<template> Specify one of the following agent templates:

Class ME Variables
Default Traffic
Failed Flow Ratio
High Bandwidth Host
High Bandwidth New App
Host Info Variables
Link ME Variables
Memory Allocation
New Application
NFPM Failed Flows
NFPM Side Unknown
Partition ME Variables
Partition Utilization
System Load
Traffic Performance
Unit Limits

Example:

agent new testagent "Class ME Variables"

agent new "agent two" "Class ME Variables"

Command Change History

Release	Modification
11.5.1	Command introduced

agent off

Disable an existing adaptive response agent without deleting it. The agent will no longer return values or create new reports, yet it can be reenabled at any time with the agent on command.

```
agent off <name>
```

where *<name>* is the name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example, *My Agent*.

Command Change History

Release	Modification
11.5.1	Command introduced

agent on

Enable an existing adaptive response agent that has been disabled. The agent will once again return values and create new reports.

```
agent on <name>
```

where *<name>* is the name of the agent to be turned on. If the agent name has a space, the words must be entered within quotation marks, for example *"My Agent."*

Command Change History

Release	Modification
11.5.1	Command introduced

agent override

For PolicyCenter / PacketShapers in Shared Configuration Mode

Override an adaptive response agent that a child configuration inherits from a parent configuration, so the agent may be modified on the child configuration. Inherited agents cannot be modified until they are overridden.

```
agent override <name>
```

where *<name>* is the name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example, *My Agent.*

agent parm

Specify the parameter values for an adaptive response agent. The agent must have been already defined with the agent new command.

```
agent parm <name> [<parm-name> <parm-value> | default]
```

<name>	Name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example, *My Agent.* If the agent name is a single word, the quotation marks are not necessary.
<parm-name>	
<parm-value>	The parameter value for *<parm-value>*, or enter **default** for the parameter s default value. For information on the acceptable and default parameter values, see the links above.
default	Return the agent to its default values

Examples:

The first example shown below changes the *ClassName* parameter for the agent *testagent* so that agent will now monitor the class */Inbound/Citrix*.

agent parm testagent ClassName /Inbound/Citrix

If you don't specify any parameters, the **agent parm** command shows current and default parameter settings for the specified agent.

agent parm "System Load"
Score Parms
RedThreshold 95(Default: 95)
GreenThreshold 90(Default: 90)

Command Change History

Release	Modification
11.5.1	Command introduced

agent show

Show data for one or many adaptive response agents, including information on the agent type and category, the corresponding plug-in file, any incident report files, and the agent version number.

The **agent show** command will show values with a timestamp based upon the end of the evaluation interval. This is different from the **measure dump** command, which shows values with a timestamp that reflects the beginning of the time interval.

```
agent show [name <name> | templates | [result <score-result>] | [feedback [<unitSN> <name>]]
```

<name>	Name of the agent. If the agent name has a space, the words must be entered within quotation marks, for example *My Agent.* If the agent name is a single word, the quotation marks are not necessary.
templates	Show a list of available adaptive response agent templates
result *<score-result>*	If the agent is unable to measure its target, the output of the **agent show name *"<name>"*** command will display additional *Result* category of data showing an explanation of the error and the error code. You can determine the meaning of an error code by issuing the command **agent show result *<score-result>*.** See the example below.
<unitSN> *<name>*	Used with the **agent show feedback** command, the ***<unitSN>*** parameter is serial number of the unit for which want to view agent data and feedback. This parameter is optional–if you do not specify a unit, the **agent show feedback** command will show data for all agents (and when issued from PolicyCenter, all agents for all configurations). The ***<name>*** parameter is the name of the agent. If no data exists for a new agent, or there is no agent with the specified name, this command will return the output No feedback available.

Examples:

The **agent show** command displays agent information for an individual PacketShaper, or when issued from the PolicyCenter client, agent information for the configuration you are editing. This information includes data on whether or not the agent has been enabled, the name of the agent, and the last score information.

For PolicyCenter configurations, an **I** to the left of the agent name indicates that the configuration has inherited that agent from a parent configuration. An **O** to the left of the agent name indicates that the configuration has a local override of an agent that supersedes the agent it inherits from its parent. An exclamation point (!) beside the agent name indicates a configuration error. Last Score Information includes the latest value measured by the agent, its status color, and the time and date of the measurement.

agent show

```
  Status            On

  Agent Name                       Status       Last Score Information
-------------------------------------------------------------------------------
  Class ME Variables agent          On          0 Yellow Wed Jan 12 02:03:00 2005 PST
  High Bandwidth New App            On          New score value in 51m 13s.
  Inbound Default Traffic           On          1 Green Wed Jan 12 02:03:00 2005 PST
  Outbound Default Traffic          On          0 Green Wed Jan 12 02:03:00 2005 PST
  Partition Utilization agent       On          0 Green Wed Jan 12 02:03:00 2005 PST
  Spoofing - Client                 On          0 Green Wed Jan 12 02:03:00 2005 PST
  Spoofing - Server                 On          0 Green Wed Jan 12 02:03:00 2005 PST
  Syn Attack - Failed Flows         On          0 Green Wed Jan 12 02:03:00 2005 PST
  Traffic Performance agent         On          1 Red Wed Jan 12 02:03:00 2005 PST
*NT = No template found for agent.
*NF = Either an action or incident file not found.
```

agent show templates

```
PacketShaper# agent show templates
                             PlugIn      Incident     Ver
Template Name                 File      Report File   Num  Category
-------------------------------------------------------------------------------
Quota Bandwidth Host            -        hostquot.cmd 1.0  Hosts
Host Info Variables             -        hostvar.cmd  1.0  Hosts
Failed Flow Ratio               -        ffratio.cmd  1.0  Hosts
NFPM Failed Flow                -        syn.cmd      1.0  Hosts
NFPM Side Unknown               -        spoof.cmd    1.0  Hosts
Link ME Variables               -        melink.cmd   1.0  User Event Emulation
Partition ME Variables          -        meptn.cmd    1.0  User Event Emulation
Class ME Variables              -        meclass.cmd  1.0  User Event Emulation
High Bandwidth New App          -        susapp.cmd   2.0  Application Health
New Application                 -        newapp.cmd   1.0  Application Health
Default Traffic                 -        dflttraf.cmd 2.0  Application Health
High Bandwidth Host             -        sushost.cmd  3.0  Hosts
Traffic Performance             -        trafperf.cmd 2.0  Network Health
Partition Utilization           -        ptnutl.cmd   3.0  Network Health
Memory Allocation               -        sysmem.cmd   1.0  Unit Health
System Load                     -        sysload.cmd  1.0  Unit Health
Unit Limits                     -        syslimit.cmd 1.0  Unit Health
```

> **Note:** The incident report files described in the above output above are the files used by each agent to create incident reports. Incident report files are different from action files, as they are used only to generate drill-down incident reports. *Do not edit or modify incident report files in any way.* Any modifications to an agent s incident report file could stop new reports from being generated for that agent.

agent show name "inbound default traffic"

```
  Agent Name          Inbound Default Traffic
  Status              On
  Template Info
    Template Name     Default Traffic
    VerNum            2.0
    Category          Application Health
    Description       This agent monitors the rate (avg-bps) of the default traffic
                      class. This agent can alert you when the amount of traffic not
                      classified (falling into 'default') is too great. This agent
                      must be used with a 'default' (i.e., /Inbound/Default) traffic
                      class.
                      Threshold Units: % of bandwidth on the partition
                      Action File Variables: $class-id, $avg-bps
    Plugin File            -
    Incident File     9.258/agent/cmd/dflttraf.cmd
    MultiInstance     Allowed
    Interval          1 minute(s)
  Score Parms
    RedThreshold       15(Default: 15)
    GreenThreshold     7(Default: 7)
    ClassName          /Inbound/default(Default: /Inbound/default)
  Color Mappings
    Green              Score < 7
    Red                Score > 15
  Last Score Status
    Value             1
    Color             Green
    Start time        Wed Jan 12 02:05:00 2005 PST
    Finish time       Wed Jan 12 02:06:00 2005 PST
  New score value in 39s.
```

If the agent in the example above had a status color of blue, the *Last Score Status* category would display additional *Result* information with an explanation of the error and an error code. The example below shows the *Last Score Status* displayingthisadditional *Result* output.

Last Score Status	
Value	0
Color	Blue
Result	Agent score parm not found. (score-result: 4569) <--------
Start time	Mon Jun 19 08:19:00 2005 PST
Finish time	Mon Jun 19 08:20:00 2005 PST

New score value in 45s.

You can determine the meaning of a *Result* error code with the **agent show result *<score-result>*** command. The following example displays information for error code 4569.

agent show result 4569
Agent score parm not found.

This next example shows the resultant output when the command **agent show feedback** is issued for a PolicyCenter configuration. (If this command was issued for a unit configuration, it will show only the agents on that individual unit.) The Feedback Information includes the latest value measured by the agent, its status color, and the time and date of the measurement.

agent show feedback

Unit	Agent Name	Feedback Information
065-10000193	ClassMeVar	1 Green Mon Jul 19 22:33:01 2005 LST
065-10000193	Hosts	1 Green Mon Jul 19 22:38:01 2005 LST
065-10000179	ClassMeVar	2 Yellow Mon Jul 19 22:33:01 2005 LST
065-10000179	Hosts	7 Yellow Mon Jul 19 22:38:01 2005 LST
065-10000238	PacketDrops	0 Green Mon Jul 19 23:00:06 2005 LST
065-10000238	Hosts	1 Green Mon Jul 19 23:00:06 2005 LST

Issue the **agent show feedback** command with the *<unit#>* and *<name>* parameters to display data for one agent on a single unit.

agent show feedback 025-10000210 "FTP Partition Over Limit"

```
Score Feedback:
  Score           18073
  Color           Red
  Category        User Event Emulation
  Start Time      Fri Oct 15 08:56:02 2005 PDT
  Finish Time     Fri Oct 15 09:56:02 2005 PDT

Incident Report Feedback:
  File Output     9.258/agent/cmd/complete/155646.htm
  Result          Success.
  Finish Time     Fri Oct 15 09:56:02 2005 PDT
```

Command Change History

Release	Modification
11.5.1	Command introduced

arp

Display and update the Address Resolution Protocol (ARP) table. The ARP table usually does not require user intervention, because it is built automatically by the ARP protocol. If you are reconfiguring or troubleshooting a network problem, you may want to manipulate the table using the **arp** command.

```
arp show|drop|flush|privadd
```

show
: Display the ARP table

drop *<ipaddress>*
: Drop the specified IP address entry from the ARP table

flush
: Flush ARP table

privadd *<ipaddress>* *<mac_addr>* *<device>*
: Add static ARP entry.

  Depending on which NICs are installed, *<device>* can be one of the following:

  Dual-port NICs:

  - slot1_in
  - slot1_out
  - slot2_in
  - slot2_out
  - slot3_in
  - slot3_out
  - slot4_in
  - slot4_out

  Quad-port NICs:

  - slot3_in1
  - slot3_out1
  - slot3_in2
  - slot3_out2
  - slot4_in1
  - slot4_out1
  - slot4_in2
  - slot4_out2

  This command will add and correlate the specified MAC address to the IP address without trying to resolve it. This will be a permanent entry in the ARP table and will be deleted only when the PacketShaper is reset. It can also be deleted manually using the **arp drop** command. If you want the entry to remain in the ARP table after a PacketShaper reset, add this command to the startup.cmd file.

Only the **show** option can be executed in look mode. All other **arp** options require touch access.

Command Change History

Release	Modification
11.3.1	Command introduced, with support for the S-series NICs

banner show

Display the messages (such as "Packet shaping: off") that are initially shown after logging into PacketShaper. You can use the **banner show** command to display all of the unit's configuration errors, warning messages, and notices. (This same information is displayed in the Info tab of the browser interface.)

```
banner show [verbose]
```

The **verbose** option displays additional information, such as the date and time and the type of message (*notice, warn*, etc.).

For example:

```
banner show
```

Packet shaping: off.

Power supply 1 FAILED.

INSIDE interface down

cat

Display the contents of a file.

```
cat [-n] <filename>
```

where **-n** numbers the output lines.

Command Change History

Release	Modification
11.2.1.8	New and improved version of the command, with an option for numbering the output lines

cd

Change your current directory.

```
cd <dir>
```

For example, type **cd 9.258/** to change to the 9.258 folder.

class capture-ids

Creates a text file named classids.txt that contains a list of all well-known class identification values. The classids.txt file is located in the 9.256/log folder. For example, the well-known ID for /Inbound is 1 and the ID for /Outbound is 2.

```
class capture-ids
```

If you use the more command to view the contents of this file, a list appears with the class ID next to each class name. This list includes all classes that can be auto-discovered not just the ones currently in the traffic tree. Part of the ID list appears below.

1 /Inbound
2 /Outbound
3 /Inbound/Inside
4 /Inbound/Outside
5 /Inbound/Default
6 /Inbound/Global
7 /Inbound/Global/IP
8 /Inbound/Global/TCP
9 /Inbound/Global/UDP
10 /Inbound/Global/Miscellaneous
11 /Inbound/Global/DECnet
12 /Inbound/Localhost
13 /Inbound/SameSide
15 /Inbound/OutsideVPNTunnel
50 /Outbound/Inside
51 /Outbound/Outside
52 /Outbound/Default
53 /Outbound/Global
54 /Outbound/Global/IP
55 /Outbound/Global/TCP

class category

Assign a traffic class to a host accounting category. (See host accounting categories for details on creating the categories.) Once you have assigned a class to a category, the bytes sent and received for the class will get tallied into the assigned category for both the source and destination hosts.

```
class category <tclass> none|<category-name>
```

You can assign multiple classes to each category, if you like. The *<tclass>* must be a leaf class; that is, you cannot assign a category to a class that has any child classes.

> **Note:** You cannot create a child class after the parent has been assigned to a host accounting category.

To remove a traffic class from a host accounting category, use:

```
class category <tclass> none
```

class copy

Copy a traffic class and its children to another parent in the traffic tree.

```
class copy <tclass><new_parent> [children]
```

Specify the explicit path and class name for the traffic class to be copied and the receiving parent traffic class. For example:

```
class copy /inbound/HTTP/Gifs /inbound/HTTP/Graphics
```

> **Note:** Any defined top talkers and top listeners are not copied with a traffic class.

class criteria

Certain services, such as Citrix and Oracle, can be further classified by application-specific criteria. For example, you can create a traffic class for a specific Citrix application or an Oracle database. You can use the **class criteria** commands to display the attributes that can be specified in a matching rule for these applications and to discover the values that can be specified for the attributes.

class criteria attributes|recent|track

attributes	Display the available application-specific criteria
recent	Show recently tracked criteria values for a class
track	Enable or disable criteria tracking for a class

The application-specific criteria format in a matching rule is:

<application>:<attribute>:<value>

where *<application>* and *<attribute>* are as described in the table below, and *<value>* is specific to your configuration and

classification requirements.

This table shows available applications, attributes, and sample values.

Application	Service Type	Attribute	Example of Value
Citrix	Citrix-ICA	application client priority	PeopleSoft pat-pc 0
DCOM	DCOM	UUID	1cbcad78-df0b-4934-b558-87839ea501c9
DICOM	DICOM	server client	DICOM_STORAGE DICOM_ECHO
FTP	FTP-Data-Clear	FileName	*.mp3
Web	HTTP	host url content-type user-agent	207.78.98.18 /Images/*.jpeg image/gif Mozilla/4.0
HTTP-Tunnel	HTTP-Tunnel	host port	207.78.98.18 80
ICMP	ICMP	type	echo
NNTP	NNTP-Clear	GroupName	alt.binaries.* microsoft.public.games
Oracle	Oracle-netv2	dbname	corp
PostgreSQL	PostgreSQL	dbname	corp
RTCP	RTCP-I	encoding media clock	GSM v 8000

Application	Service Type	Attribute	Example of Value
RTP	RTP-I	encoding media clock to from user-agent source destination	dynamic, GSM v 8000 +12125551212@server.phone2net.com agb@bell-telephone.com *X*Lite* 207.78.98.18 207.78.98.18
			Additional Information ■ For SIP attributes, you can enter a substring of the attribute. For example, to match all Motorola models, you can enter **Motorola** for the user-agent criteria. ■ The asterisk (*) wildcard is supported for user-agent.
SMTP	SMTP-Clear	SenderEmail	*@bluecoat.com
SOAP	SOAP-HTTP	host url content-type user-agent	207.78.98.18 /Images/*.jpeg image/gif Mozilla/4.0
SSL	SSL	common name	optionslink.etrade.com
WAP	WAP	WAPURI	*.bluecoat.com

> **Note:** In order to add an application-specific matching rule to a class, the class' service type must be the one indicated in the Service Type column above. For example, to classify by Oracle database name, the class must be based on the Oracle-netv2 service.

You can use the **class criteria** commands to identify the specific values to use in application-specific matching rules. First, you use class criteria attributes to get a list of applications and attributes that can be used in matching rules. Next, you use class criteria track to enable tracking on a specific class. Then, you use class criteria recent to see a list of recent values for the class; the output will provide you with the information you need to create an application-specific matching rule.

The following example shows how you can use the **class criteria track** and **class criteria recent** commands to identify the SIP user-agent for RTP-I traffic:

```
PacketShaper# class criteria track /inbound/RTP-I/Default RTP user-agent
```

After a period of time in which VoIP calls are made, issue the following command:

```
PacketShaper# class criteria recent inbound/RTP-I/Default
       Traffic Class: /Inbound/RTP-I/Default
```

```
       Application: RTP
       Attribute: user-agent (SIP User-Agent)
Recent Attribute Values (most recent first)
------------------------------------------------------------------------------
   1. *X*Lite*
```

The above output indicates that **X*Lite** is the string that should be specified as the SIP user-agent criterion.

Command Change History

Release	Modification
	UUID attribute for DCOM added
8.2.0	dynamic encoding attribute for RTP added (for classification of dynamic codec numbers)
	Wilcard (*) support added for RTP-I user-agent attribute.

class criteria attributes

Display a list of applications that can be further classified using application-specific criteria. This information can be used to create a matching rule for a traffic class.

class criteria attributes

```
Application   Attribute       Example of Value
------------------------------------------------------------------------------
Citrix        application     PeopleSoft
              client          pat-pc
              priority        3
DCOM          UUID            1cbcad78-df0b-4934-b558-87839ea501c9
DICOM         server          DICOM_STORAGE
              client          DICOM_ECHO
FTP           FileName        *.mp3
Web           host            207.78.98.18
              url             /Images/*.jpeg
              content-type    image/gif
              user-agent      Mozilla/4.0
HTTP-Tunnel   host            207.78.98.18
              port            80
ICMP          type            echo
NNTP          GroupName       alt.binaries.*
Oracle        dbname          corp
PostgreSQL    dbname          corp
ProxySG       app/port        HTTP/6000
Identity
RTCP          encoding        GSM
              media           v
```

```
              clock           8000
RTP           encoding        GSM
              media           v
              clock           8000
              to              +12125551212@server.phone2net.com
              from            agb@bell-telephone.com
              user-agent      Motorola VT1000
              source          207.78.98.18
              destination     207.78.98.18
SMTP          SenderEmail     *@packeteer.com
SOAP          host            207.78.98.18
              url             /Images/*.jpeg
              content-type    image/gif
              user-agent      Mozilla/4.0
SSL           commonName      www.packeteer.com
WAP           WAPURI          *.packeteer.com
```

The application-specific criteria format in a matching rule is:

<application>:<attribute>:<value>

where *<application>* and *<attribute>* are as described in the table above, and *<value>* is specific to your configuration and classification requirements. (The values can be determined with the class criteria track command.) For example, the criterion for creating a matching rule that matches traffic to the PeopleSoft application running over Citrix is:

citrix:application:PeopleSoft

Examples of other valid matching rules are:

ftp:FileName:*.mp3

icmp:type:echo

web:host:207.78.98.18

oracle:dbname:corp

> **Note:** These matching rules can be specified in the class new or class rule commands.

Command Change History

Release	Modification
8.2.0	UUID attribute for DCOM added dynamic encoding attribute for RTP added (for classification of dynamic codec numbers) Wilcard (*) support added for RTP-I user-agent attribute.

class criteria recent

Display a list of values that have been tracked for a specific traffic class. This command is used after criteria tracking has been enabled.

class criteria recent *<tclass>*

where *<tclass>* is the traffic class that has criteria tracking enabled with the class criteria track command.

Example:

class criteria recent /inbound/http

Traffic Class: /Inbound/HTTP
Application: Web
Attribute: content-type (Content Type)

Recent Attribute Values (most recent first)
---

1. text/html
2. image/gif
3. application/x-shockwave-flash
4. image/jpeg
5. application/x-javascript
6. application/x-www-form-urlencoded
7. text/css
8. application/octet-stream
9. text/plain
10. text/xml

class criteria track

Enable application-specific criteria tracking on a traffic class. (Use class criteria attributes to see a list of valid applications and attributes.)

class criteria track *<tclass>* *<application>*|off *<attribute>*

where *<tclass>* is the traffic class, *<application>* is the application name, and *<attribute>* is the type of values you want to track. For example, to set up criteria tracking of content-types for the HTTP class:

class criteria track /inbound/HTTP web content-type

After you have enabled criteria tracking for a class, use the class criteria recent command to see a list of recent values for the class.

When you are done, turn off tracking using the following command:

class criteria track *<tclass>* off

For example, to turn off criteria tracking on the /inbound/http class:

class criteria track /inbound/http off

class delete

Remove a class from the traffic tree.

```
class delete <tclass> [children]
```

<tclass>	The name of the traffic class to delete. The class' explicit hierarchical path must be supplied only if the class name itself is not unique.
[children]	Specify to delete all of the class' child classes; this parameter is required in order to delete a class that has children

If you delete a class that was created by traffic discovery and you have traffic discovery turned on, the class is likely to appear again in your traffic class tree.

class discover

Enable or disable traffic discovery within a specific class. For class discovery to take effect, traffic discovery must be enabled at a global level using the **setup discover** command.

```
class discover <tclass> [inside|outside|both|off]
```

<tclass>	The name of the traffic class within which you are enabling or disabling traffic discovery. The class' explicit hierarchical path must be supplied only if the class name itself is not unique.
[inside|outside|both|off]	Specify the location of the server for which you want traffic to be discovered, or off to turn off discovery for this class. If you don't specify one of these options, the action defaults to turning on traffic discovery effectively using the **both** setting.

class group delete

Delete a custom service group. When you delete a group, the services in that group are moved into the Unassigned group. Note that you cannot delete the built-in groups.

```
class group delete <group_name>
```

where *<group_name>* is the name of the custom group.

Example:

```
class group delete CorpApps
```

- The child configuration inherits the parent's group of the same name.
- This group will contain all the services defined in the inherited group except for any services that were moved out of the group when the group had been overridden. These services will stay in the group that they had been moved into.
- If the overrriden group contained other services that were moved into it, these services will go into Unassigned after the group is deleted.

Command Change History

Release	Modification
8.5.1	**class group** command introduced

class group move

Move all services from one service group into another, or move one service into a different group. You can move services into a built-in or custom group.

```
class group move {<group_name> | <group_name>:<service_name>} <group_name>
```

Examples:

To move all the services in the Multimedia group into Mygroup:

```
class group move multimedia mygroup
```

To move service Citrix from the RemoteAccess group into the ClientServer group:

```
class group move RemoteAccess:citrix ClientServer
```

Additional Information

- To move a service back into its default group, use the class group reset command.
- Services in the NonIPv4 group cannot be moved to other groups.

Command Change History

Release	Modification
8.5.1	**class group** command introduced

class group new

Create a custom service group. PacketShaper includes 25 built-in service groups, but if these don't suit your needs, you can create your own groups. For example, if you have created user-defined services for your custom applications, you may want to create a custom group for them. You can create up to 25 custom service groups.

```
class group new <group_name> <description>
```

where *<group_name>* can be up to 31 characters (including hyphens, underscores, and periods) and *<description>* can contain up to 80 characters. If the description contains spaces, you must enclose the text string in quotes.

Example:

```
class group new CorpApps "Corporate applications"
```

Command Change History

Release	Modification
8.5.1	**class group** command introduced

class group override

For PolicyCenter / PacketShapers in Shared Configuration Mode

Override a service group that a child configuration has inherited from a parent configuration. Use this command if you don't want the service group to inherit any more changes from the parent.

```
class group override <group_name>
```

where *<group_name>* is the name of the built-in or custom service group that has been inherited.

> **Note:** After overriding a service group, if you then want to re-inherit it from the parent configuration, you can delete the overridden group. Or, you can re-inherit *all* service groups from the parent with the class group reinherit all command.

class group reinherit

For PolicyCenter / PacketShapers in Shared Configuration Mode

Delete all service groups from the current configuration and re-inherit the service groups from the parent configuration. Use this command when a child configuration contains a number of local overridden service groups and you decide that you want the configuration to go back to inheriting the parent's groups.

```
class group reinherit all
```

Additional Information

- You may decide to perform this operation if, after modifying service groups in a child configuration, you end up with configuration errors (service conflicts) that you can't resolve.
- Use the class group show command to confirm that each service group has the **I** (Inherited) marker, indicating the group is inherited from the parent configuration. Any local custom groups you had in the child configuration will no longer appear on the group list.

class group reset

Return all services back to their default groups or return a single service back to its default group. Use this command if you have moved services around to different groups and then discover you made a mistake or have changed your mind.

```
class group reset <service_name> | all
```

where *<service_name>* is the name of the service that you want to return to its default built-in group.

Examples:

To move the Citrix service back into its original, default group:

```
class group reset citrix
Service citrix successfully reset to its default group (RemoteAccess).
```

To return all moved services back to their default built-in groups:

```
class group reset all
All services reset to their default group.
```

> **Note:** If you had created any custom groups, these groups will remain after the **reset all**, although they will no longer contain any services.

Command Change History

Release	Modification
8.5.1	**class group** command introduced

class group show

Display a list of all group names and their descriptions, details about a particular group, or a list of all services and the group to which each belongs. You can also use this command to find out to which group a particular service belongs.

```
class group show [<group_name> | service | {service <service_name>}]
```

Examples:

To display a list of all groups and their descriptions:

```
class group show
```

To list all groups and all the services that belong to each group:

```
class group show service
```

To display details for a particular group:

```
class group show healthcare


Name         : Healthcare
Description  : Healthcare related applications
Num services : 2
Services in Healthcare group
        DICOM                       Digital Imaging and Communications in Me
        HL7                         Health Level Seven (HL7)
```

To find out to which group a service belongs:

```
class group show service ftp
  service:"ftp" belongs to group:"Internet"
```

Command Change History

Release	Modification
8.5.1	**class group** command introduced

class hosts

Displays a list of all host references in matching rules and host lists. A host may be listed as an IP address, a DNS name, or an LDAP DN (Lightweight Directory Access Protocol domain name) for a host list entry. If more than one matching rule contains the same host reference, the host is shown only once.

```
class hosts
```

Host reference	DNS name	IP address
127.0.0.3	-	127.0.0.3
	www.lycos.com	206.79.171.51...
	www.excite.com	198.3.98.99

If a DNS name resolves to more than one address, the first address is listed followed by an ellipsis (...). To list the additional addresses, use the **dns lookup** command.

If there is a problem resolving a DNS name, the third column shows the DNS error message.

class id

Change or view a traffic class identification number. The numeric ID of a class is used for Simple Network Management Protocol (SNMP) and the measurement engine. It must be unique and does not change when the class is renamed.

```
class id <tclass> [<number>]
```

<tclass>	The name of the traffic class whose ID you are changing. The class' explicit hierarchical path must be supplied only if the class name itself is not unique.
[*<number>*]	The new unique number for the traffic class

> **Note:** Class IDs should be changed in special circumstances only, for example when you want class IDs to be the same across multiple PacketShapers. Changing class IDs can lead to erroneous reporting of data if you choose an ID value that was previously used by another class.

To see the current ID for a traffic class, type **class id *<tclass>***. To see the ID for all services, use the class services id command.

class licenses

Limit the number of TCP flows allowed simultaneously in the given class, where the number of flows admitted to a class is based on a fixed number instead of the available bandwidth.

```
class licenses <tclass> off|<number>
```

where *<number>* is the maximum number of TCP flows to admit.

After *<number>* flows are active on the specified traffic class, new flows are given the admission control treatment defined by **policy admit**.

Specify the **off** option to remove the limit on the number of flows.

After you have limited the flows with the **class licenses** command, you can use the **traffic licenses** command to see the number of flows currently in use.

class load

Load a new traffic configuration file. This command will load the traffic tree and everything related to the classes in the tree, such as policies and partitions. This feature can be used to share configurations with other units. You can upload a saved configuration to the 9.256/ folder of another PacketShaper unit and then activate it with the **class load** command.

```
class load <path>
```

The following example loads a config.ldi file from the 9.256/ folder:

```
class load 9.256/config.ldi
```

The **class load** command prompts for confirmation, then overwrites the existing 9.256/cfg/config.ldi file with the file you specify.

class move

Relocate a traffic class by assigning it to a new parent class. Unlike using the **class copy** command, the class will no longer reside under its original parent, but will be moved to a new location in the tree structure.

```
class move <tclass><new parent> [children]
```

Use the literal **children** to move all of the class' children as well; otherwise, only the parent class will be moved and the children will be promoted a level.

> **Note:** When moving a traffic class, you cannot change the direction. For example, you cannot move a traffic class from /Inbound to /Outbound.

class new

Create a new traffic class.

```
class new <parent_name><name> [nodefault]  <rule>
```

<parent_name>	The parent class for the new traffic class. You must use the explicit hierarchical pathname if the class name is not unique - for example, */inbound/http*.
<name>	A unique name for the new traffic class, up to 31 characters long. Use only alphanumeric characters and the following special characters: underscore (_), hyphen (-), and period (.). Specify only the class name, without the leading tree hierarchy pathname.
[nodefault]	A Default match-all class will not be created (applicable when creating a child class). For example, if you don't specify the **nodefault** parameter when creating the Inbound/HTTP/WebSurfing class, PacketShaper will also create an Inbound/HTTP/Default class. If you do specify the **nodefault** parameter, the Inbound/HTTP/Default class will *not* be created.
<rule>	A matching rule defines a traffic class' attributes. A class can contain multiple matching rules, which are treated as separate, distinct rules. To define one or more rules for a traffic class, see **class rule**. For matching rule details, see Matching Rule Details in PacketGuide.

You cannot create a child class if the parent has been assigned a host accounting category.

> **Note:** If your unit is within one traffic class of its capacity, PacketShaper will not let you create any more classes. This is due to the possibility that two classes will be created in some circumstances. For example, when you create the first child class for a parent, a Default class automatically gets created.

Creating a Class for a Specific File Type

Specify GIF file downloads:

```
class new inbound/http graphics outside service:http web:url:"*.gif"
```

Specify MP3 files downloaded via FTP:

```
class new inbound/ftp ftp_mp3_downloads outside service:ftp-data-clear ftp:filename:*.mp3
```

Creating a Class for a Specific Host or Port

Target any traffic from an external host:

```
class new inbound competitor outside host:145.34.0.2 service:http
```

Specify web traffic to a port other than port 80, the normal web port:

```
class new inbound web_in inside service:http port:8080
```

Creating a Class for a Specific URL, IP Address, or Host List

Specify a URL (http://altman.com/support/support.htm):

```
class new inbound altman outside service:http host:altman.com
web:url:"/support/support.htm"
```

For security purposes, you can classify TCP traffic based on the origin of the connection. To do this, create a traffic class that specifies an outside TCP client. Create this type of class only after you are satisfied that traffic discovery has sufficiently identified traffic on your network. Otherwise, it will prevent the discovery of more specific services.

```
class new inbound mystuff outside tcp client
```

Specify an IP address if you do not have a DNS server configured:

```
class new inbound server_guru inside 203.160.106.3
```

Specify a host list (a set of IP addresses and/or DNS names):

```
class new inbound/servers inside host:any outside list:servers
```

Creating a Class for an IPv6 Subnet

```
class new inbound ipv6-2 inside net:2001:db8:1234:5678::/64
```

Creating a Class for ICMP or IGMP Traffic

When creating a symmetrical traffic class for the ICMP or IGMP protocols, we recommend that you explicitly specify the protocol for both the inside and outside interface. For example:

```
class new /Inbound/ABQ ICMP inside ICMP outside ICMP
```

To create an asymmetrical traffic class for ICMP or IGMP, where traffic is classified on either the inside or the outside interface:

```
 class new /Inbound/ABQ ICMP inside ICMP
```

or

```
class new /Inbound/ABQ ICMP outside ICMP
```

class note

Annotate a traffic class.

```
class note <tclass>"<note>"
```

This note appears in the **class show** display. Non-printing characters are not allowed.

class owner

Specify an owner name for a traffic class.

```
class owner <tclass> <ownername>
```

The owner name can be up to 32 characters and the following special characters are not allowed: quote ( ), ampersand (&), backslash (\), and non-printing characters.

The owner name appears in the **class show** display.

class rename

Rename a traffic class.

```
class rename <tclass><new tclass>
```

The class to be renamed must be specified with its full pathname; do not specify the path for the new class name. (The path from the original name is used.) For example:

```
class rename inbound/test sap
```

When renaming a class you are not allowed to change just the case; for example, you cannot rename HTTP to http.

> **Note:** If you rename a class and that class has an event associated with it, the class name is not automatically updated in the event registration. Therefore, after renaming a class, you will need to re-register the event with the new class name.

class reset

Clears all classes, policies, and partitions and reverts to either the default or model tree.

```
class reset [model]
```

The **model** option resets the configuration to a pre-configured tree that can be used as is, or modified to suit your needs. This tree organizes network traffic into folders of common categories, such as VoIP, risky websites, business-critical applications and data, and recreational websites and applications. It includes classes based on service groups or URL categories.

If you reset the tree without the **model** option, the tree is reset to the default—a bare-bones traffic tree that includes /Inbound and /Outbound classes with a Default class for each, and a Localhost class for the inbound and outbound directions. You can build out this tree by turning on traffic discovery or by manually creating classes for the type of traffic you want to track.

Notes:

- Use the config save CLI command to back up your configuration before resetting the tree; this gives you the capability of restoring the traffic tree and configuration if you change your mind.

- If a number of classes have top talkers/listeners enabled, a class reset could take several minutes when many hosts are being tracked. Before issuing the class reset command, you may want to disable top talkers/listeners to avoid this delay. See hostdb topusers.

Command Change History

Release	Modification
8.7.1	**model** option introduced

class rule

Add or delete matching rules.

```
class rule add <tclass> <rule>

class rule delete <tclass> <rule_id>
```

The maximum number of matching rules per traffic class depends on the PacketShaper model. If a traffic class has more than one matching rule, PacketShaper compares the flow to the first specification. If it doesn't find a match, it moves to the class' next matching rule.

Matching rules are identified by a rule ID in brackets [ ]. You can determine the rule ID by using the command: **class show *<tclass>***

See Matching Rule Details in PacketGuide for additional information.

Examples:

Create a new Oracle class with three matching rules. The first matches on an inside host IP address of 190.160.0.207, the second matches on 190.160.0.208, and the third on 190.169.0.254.

```
class new /outbound oracle inside service:oracle host:190.160.0.207
 class rule add /outbound/oracle inside service:oracle host:190.160.0.208
 class rule add /outbound/oracle inside service:oracle host:190.169.0.254
```

Create a new FTP class with two matching rules, one for the outside and the other for the inside.

```
class new inbound/ftp ftp_mp3_downloads outside service:ftp-data-clear ftp:filename:*.mp3
 class rule add inbound/ftp/ftp_mp3_downloads inside service:ftp-data-clear
ftp:filename:*.mp3
```

Recall that if a traffic class has more than one matching rule definition, PacketShaper compares the flow to the first specification. If it doesn't find a match, it moves to the class' next rule. Traffic that matches any of a class' matching rules will fall into the class.

If the info page has flagged one or more of your classes with the configuration error message *attrib iqosMatchingRule = ??? , Failed to add matching rule to traffic class*, you have exceeded the maximum number of matching rules available on your PacketShaper model. (In the CLI, you can display configuration error messages with the **class show *<tclass>*** command.) To free up resources, you need to remove one or more classes or matching rules. Configuration errors will disappear once the total number of matching rules is less than the unit s limit. If you find that you are consistently exceeding your unit s maximum configuration limits, you should consider upgrading your PacketShaper.

class services

List the services available in PacketShaper.

```
class services [<service name>] [id]
```

<service name>	The name of a service; you can type the complete name, or just the first few letters
[id]	List the internal ID numbers associated with each service name. Service ID numbers are recorded in flow detail records (FDRs). Since FDRs record the service ID, not the service name, the **class services id** command would be useful for someone interpreting FDR data with a protocol analyzer or other tool that displays FDR data.

The *<service name>* option is useful for narrowing down the service list to a particular name you are looking for. The following example lists all the services that start with AOL:

```
class services aol
```

```
AOL-IM AOL - Instant Messenger & ICQ Client-Server
 AOL-IM-File AOL-IM - Point to Point File Transfer
 AOL-IM-ICQ AOL - Instant Messenger & ICQ2000
 AOL-IM-IMAGE AOL-IM-Image - Point to Point Chat
 AOL-IM-Talk AOL-IM - Point to Point Talk
```

Command Change History

Release	Modification
11.1.1	Removed [plug-in] option

class set

Make a traffic class an exception class, or configure a class to allow its policy to be inheritable.

```
class set <tclass> inherit|standard|exception
```

inherit	Inheritable traffic classes have policies that can be applied to other classes when the other class doesn't have its own policy. Specific rules apply to how PacketShaper decides which policy a class should inherit; see Inheritance Rules in PacketGuide for details. The output of the class show command indicates (with an I flag) which classes have an inheritable policy.
standard	Standard traffic types have no exception or inheritable attributes.
exception	Exception traffic classes are always positioned above non-exception classes in the tree. When you make a class an exception class, you redefine the search order that PacketShaper uses to find a match for traffic flow. The exception attribute can be applied to all classes except /Inbound, /Outbound, and any default match-all classes. Marking a traffic class as an exception ensures that it is ordered first in the subtree, overriding the tree's built-in hierarchical order.

class show

Display traffic class information for a specific class, standby partner, or the entire traffic tree.

`class show [partner <partner_ip>] [<tclass> | verbose <tclass> | since <seconds>] | [id]`

where

partner *<partner_ip>*	Lists the traffic classes of a standby partner
<tclass>	Lists configuration details about the specified traffic class, for example matching rules and policies
verbose *<tclass>*	Lists all host lists referenced by a traffic class
since *<seconds>*	Shows only classes auto-discovered within the last number of seconds

Example:

```
class show dhcp
```

```
Traffic Class: /Outbound/DHCP
Partition: /Outbound
Class Flags: autocreated
Rule Types: optimized
Current guaranteed rate 0              excess rate 0
Matching Rules:
    [52 ]   inside      any host    service:Client          any port    UDP
             outside     any host     service:DHCP-S          any port
    [54 ]   inside      any host    service:Client          any port    UDP
             outside     any host     service:DHCP-C          any port
    [53 ]   inside      any host    service:DHCP-S          any port    UDP
             outside     any host     service:Client          any port
    [55 ]   inside      any host    service:DHCP-C          any port    UDP
             outside     any host     service:Client          any port
 no policy
 Class id (for SNMP and Measurement Engine): 1069
```

The **Class Flags** indicate class attributes:

autocreated—The class was created with the traffic discovery feature.

built-in—One of the classes built into PacketShaper (such as Inbound and Outbound). Built-in classes cannot be deleted.

cacheable—The class is cacheable (that is, a class based on an IP address that is on the same side as the cache).

discovering—Traffic discovery is turned on for this class.

exception—The class is treated as an exception, overriding PacketShaper's default ordering.

inherited—The policy for the class is inheritable.

policy—The class has a policy. (The specific policy type is shown next to Policy Flags near the bottom of the output.)

The **Rule Types** indicate the type of matching rule:

optimized—The class is optimized. An optimized class is one that was auto-discovered or one that was manually created with a simple matching rule (service type, IP address, or port number).

address-is-cacheable—The class has a pure IP address-based matching rule that is on the same side as the cache (on the inside, by default). It can be an individual IP address, a range of IP addresses, an address with a mask, or host lists. These classes can be cached unless an error in the tree configuration is causing cacheability problems.

match-all—This class is a match-all class (protocol = any, service = any; for example, a Default bucket).

When you specify a standby partner, the output shows the classes in the partner's traffic tree. You can also display details about a specific class in a standby partner's traffic tree, as shown in the example below.

```
class show partner 10.9.64.155 outbound/citrix
Traffic Class: /Outbound/Citrix
Partition: /Outbound
Class Flags: autocreated
Rule Types:

Current guaranteed rate     0    excess rate     0

Matching Rules:
[1  ]   inside  any host  service:Client  any port  TCP
outside any host  service:Citrix-ICA  any port

[2  ]   inside  any host  service:Citrix-ICA  any port  TCP
outside any host  service:Client  any port

[3  ]   inside  any host  service:Citrix-SB  any port  UDP
outside any host  service:Client  any port

No policy
Class id (for SNMP and Measurement Engine): 16777380
```

If you don't specify a class or standby partner, all classes in the traffic tree are displayed, but with less detail. When displaying the entire traffic tree with the **class show** command (as shown in the following example), several flags indicate class attributes, type of matching rule..

```
class show
```

```
Derivation: (I)nherited (O)verride (U)nderride (L)ocal
 Class Flags: (A)utocreated (D)iscovering (E)xception (I)nherit (P)olicy (C)acheable
 Rule Types: (o)ptimized (m)atch-all (a)ddress is cacheable

 Class Name                      Flags    Partition Name

 Inbound                           m  /Inbound
  Localhost                      E P     /Inbound
  10.7.38.0                      a /Inbound
  CUSTOMER                       P   ma /Inbound
  mysite.org                     C a /Inbound
  Default                        IP   m  /Inbound
  Outbound                          m  /Outbound
  Localhost                      E P     /Outbound
  10.7.38.0                      a /Outbound
  CUSTOMER                       ma /Outboun
  mysite.org                     C a /Outbound
  Default                        IP   m  /Outbound
```

Command Change History

Release	Modification
11.6.1	**partner** option added

class test

Test a traffic flow against the present classification tree in order to determine the flow's class, partition, and policy.

```
class test <direction> <protocol> [<inhost:inport> <outhost:outport>] [<device>]
```

<direction>	**inbound** or **outbound**
<protocol>	**tcp**, **udp**,**icmp**, **netbeui**, **ipx**, **appletalk**, **decnet**, **fna**, **sna**, **lat**, or **misc**
<inhost:inport>	The inside and outside IPv4 or IPv6 addresses and port numbers to test (required for IP protocols only: TCP, UDP, ICMP)
<outhost:outport>	You must supply both an inside and an outside address. Use **0.0.0.0:0** as a placeholder if you don't have an address to test on one of the sides. If the hosts are IPv6 addresses, surround the IPv6 address with square brackets. For example, *[2000:1:2::1]:3456* where *3456* is the port number.
<device>	A valid device name, such as **slot1** or **slot3_pair1**

This information simulates a flow, returning the following information:

Traffic Class	The traffic class in the current traffic tree into which the flow would be classified
Partition	The partition associated with the matching traffic class. If the traffic doesn't have its own partition, the parent partition is used.
Policy	The matching policy. If the matching traffic class has no applied policy, the policy is inherited. See Policy Inheritance Rules in PacketGuide.

> **Note:** The **class test** command will only match traffic classes that have "any" for the server location.

Examples:

```
class test inbound appletalk

  Traffic class  --> /Inbound/AppleTalk
  Partition      --> /Inbound
  Policy         --> /Inbound/Default
```

```
class test inbound tcp 216.110.182.168:80 0.0.0.0:0

  Traffic class  --> /Inbound/HTTP
  Partition      --> /Inbound
  Policy         --> /Inbound/Default
```

Additional Information

- The **class test** command can be used to test basic classification for IP protocols, but is not intended to test every type of classification PacketShaper offers. Its purpose is to check a particular IP address or port number to determine how the traffic is classified into existing port-based and IP address-based classes in the traffic tree. The command does not include fields for specifying more complex types of classification such as MAC address.
- The **class test** command requires touch access.

Command Change History

Release	Modification
11.2.1	*<device>* option introduced

class user-group

List the names of user groups in a specific Active Directory domain or in all domains.

```
class user-group <domain_name>|all
```

This command is part of the user awareness feature and requires that a BCAAA server beinstalled and configured.

If you have a long list of user groups, some may scroll off the screen; if you want to be able to scroll through the list, use one of the following techniques:

- Output the list to a text file, for example: **class user-group all > grouplist**. When the command prompt redisplays, the file has finished saving. (This might take a while.) To display the list a page at a time: **more grouplist**
- Turn on session logging in your remote login utility (such as Putty or SecureCRT) before issuing the **class user-group all** command. You can then open the log file in a text editor.

Command Change History

Release	Modification
9.2.2	**class user-group** command introduced

class user-services delete

Remove a user-defined service.

```
class user-services delete <serviceName>|all
```

where *<serviceName>* is the name of the service you want to delete. Use the **all** parameter to delete all user-defined services.

Example:

```
class user-services delete TDemployees
```

Additional Information

- Service names are case sensitive. You must enter the service name with the same upper/lower case with which it was created.
- Use the class user-services show command to see a list of services that have been user-defined.

Command Change History

Release	Modification
8.4.3	**all** parameter introduced
8.4.1	**class user-services** command introduced

class user-services new

Create a custom service in order to identify and categorize traffic that is not currently classified by PacketShaper, or that is classified into a different service. This command allows you to create services for in-house applications on your network. The service can be defined by a signature (hex or string) and/or by port numbers.

```
class user-services new <serviceName> [signature:<hex>|<string> offset:<offset_value>]
[port:<nnn[-nnn]>] [packets:<packet_value>] [ipproto:TCP|UDP] [description:<string>]
```

<serviceName>	The name of the service, up to 30 characters long. Use only alphanumeric characters and the following special characters: underscore (_), hyphen (-), and period (.). The service name is case sensitive.
signature	The signature can be specified in hexadecimal format or as a quoted string. The string can be up to 30 characters long, is case sensitive, and must be enclosed in quotation marks. The hex representation can be up to 30 characters long. It must begin with 0x.
offset	Starting position of the signature in the payload (after the header). Valid values for the offset are 0-1499.
packets	Number of inbound or outbound data packets in each new flow that will be inspected for the signature. Up to 10 packets in each direction can be inspected. **Note:** Packets in each direction are counted separately. For example, a value of 8 tells the PacketShaper to look for the signature in the first eight inbound packets and first eight outbound packets of each new flow.
port	The port number or a range of port numbers. If the port option is not specified, the PacketShaper will inspect traffic on all ports.
ipproto	Type of IP protocol (UDP or TCP)
description	A description of the user-defined service, enclosed in quotation marks; up to 80 characters long.

The following types of traffic are candidates to be classified as a user-defined service:

1) traffic that PacketShaper has identified as an unknown service,

2) applications that have user-configurable ports (such as peer-to-peer and instant messaging)

Services that are associated with well-known ports (such as HTTP on port 80, FTP on port 21, and NNTP on port 119) cannot be classified into a user-defined service.

Examples:

```
class user-services new TDemployees signature:"TD Employee" offset:6 packets:1
description:"TD Employee Database"
```

```
class user-services new BCpayroll signature:0x424320706179726F6C6C offset:0 description:"BC
Payroll application"
```

Additional Information

- You can use a third-party network protocol analyzer, such as EtherPeek or Wireshark, to analyze a trace to get the signature.
- You can create up to 10 user-defined services (UDS).
- The user-defined services are auto-discoverable.
- The user-defined services are stored in the config.ldi configuration file.

Command Change History

Release	Modification
8.4.1	**class user-services** command introduced

class user-services show

Display a list of user-defined services.

```
class user-services show [<serviceName>]
```

Example:

```
class user-services show

User Defined Services
         1. Name:BCpayroll
                serviceid:645 signature:0x4243207061797 26f6c6c offset:0 packets:2
                ipproto:TCP/UDP description:"BC Payroll application"
         2. Name:TDemployees
                serviceid:647 signature:"TD Employee" offset:6 packets:1
                ipproto:TCP/UDP description:"TD Employee Database"
```

Command Change History

Release	Modification
8.4.1	**class user-services** command introduced

class users

List the names of users in a specific Active Directory domain or all domains.

```
class users <domain_name>|all
```

This command is part of the user awareness feature and requires that a BCAAA server be installed and configured.

A long list of users will scroll off the screen; if you want to be able to scroll through the list, use one of the following techniques:

- Output the list to a text file, for example: **class users all > userlist**. When the command prompt redisplays, the file has finished saving. (This might take a while.) To display the list a page at a time: **more userlist**
- Turn on session logging in your remote login utility (such as Putty or SecureCRT) before issuing the **class users all** command. You can then open the log file in a text editor.

Command Change History

Release	Modification
9.2.1	**class users** command introduced

class web-app disable

Disables the service of a web-based application (such as Facebook or YouTube). After a service is disabled, the traffic will get classified as HTTP or SSL. You might want to disable a service when you prefer to control the traffic by its URL category. For example, you can disable the Facebook service and then control all social networking traffic with a single Social Networking URL category class. This technique helps conserve classes and provides an easy way to report on and control how much HTTP is on the network.

```
class web-app disable <service>
```

where *<service>* is the name of the web application to disable. To see a list of web services that can be disabled, use the **help class web-app** command. For example, the help output includes the followin:

```
The following values are currently supported for this argument:
  Facebook
  Youtube
  MySpace
  Flickr
  GoogleVideo
  Ogg
  Smugmug
  Ofoto
  Motion
  WebShots
```

PolicyCenter Support

Web services cannot be disabled in PolicyCenter; this feature is supported in local mode only. If PolicyCenter pushes a class to a PacketShaper that is in shared mode, and that class uses a service that has been disabled locally on the PacketShaper, the class will still be created. However, traffic will not get classified into the class as long as the service is disabled.

Additional Information

- If the web browser interface is open when you enable/disable the service, you will need to refresh the browser window to load the configuration change.
- When a web application is disabled, you cannot create a class for that service in the CLI, Advanced UI, or Sky UI. The application will not appear on the **Services** drop-down list in the Sky or Advanced UIs after it has been disabled.
- Disabled web applications will not get classified into the service group to which the service belonged.
- If a class already exists for a disabled web application, the class will still appear in the traffic tree but will no longer get any class hits. It will *not* have a configuration error.

Command Change History

Release	Modification
8.7.1	**class web-app** command introduced

class web-app enable

Re-enables a web application (such as Facebook or YouTube) after it has been disabled. You would use this command if you change your mind about disabling a service and want to start classifying the service separately again.

```
class web-app enable <service>
```

where *<service>* is the name of the web application to enable. To see a list of web services that are currently disabled, use the class web-app show command.

Additional Information

- If the web browser interface is open when you enable/disable the service, you will need to refresh the browser window to load the configuration change.
- To verify that the service has been re-enabled, use the class web-app show command. The service should *not* be listed in the output of the **show** command.
- After you have re-enabled a web service, you can manually create classes based on this service or let the PacketShaper auto-discover the class. If the class already existed in the tree, it will start getting class hits once the service is re-enabled.

Command Change History

Release	Modification
8.7.1	**class web-app** command introduced

class web-app show

Lists the services that have been disabled with the class web-app disable command.

```
class web-app show
```

Example:

```
class web-app show

Classification of the following Web applications is disabled:
        Facebook
        YouTube
```

Command Change History	
Release	**Modification**
8.7.1	**class web-app** command introduced

config clear

For PolicyCenter / PacketShapers in Shared Configuration Mode

Clears all non-default configuration values from the named configuration. If none is named, it clears the current configuration. Clearing a child configuration means that the child will derive its sharable attributes and settings from its parent configuration. If you clear a parent configuration, its child configurations will no longer inherit any values from its parent.

```
config clear [<cfg_path>]
```

config dump

For PolicyCenter/ PacketShapers in Shared Configuration Mode

This command prints out the current effective configuration objects' formats and attributes in something like LDAP data interchange format. Useful mainly for development and diagnostic purposes.

```
config dump
```

config errors

Display configuration errors for the unit. When issued from PolicyCenter, this command displays errors for the PolicyCenterconfiguration currently being edited.

```
config errors
```

> **Note:** Configuration errors are also shown in the output of the banner show command.

config load

Load saved configuration files (such as config.ldi and config.cmd). Sharable settings are saved in files with the .ldi file extension, while nonsharable settings are saved in the .cmd file.

This command can load the traffic tree, partitions, policies, host lists, events, agents, basic settings (such as shaping and traffic discovery), security settings (such as passwords and login access protocols), SNMP, SNTP, and Syslog settings, DNS server, and gateway addresses, domain names, time zones, and network interface settings.

> **Note:** Use the setup show command to see a list of sharable and nonsharable settings that are stored in the configuration files.

```
config load <file> [<cfg_path>] [complete]
```

<file>	The name of a saved configuration file. Include the .ldi file extension to load just an .ldi file, or include the **complete** parameter to load both the .ldi and .cmd file associated with the specified .ldi filename. This command loads files from the PacketShaper 9.256/ folder or the PolicyCenter directory .
<cfg-path>	When you issue this command from PolicyCenter, include the path of the PolicyCenter configuration to which you want to load the file(s).
[complete]	Include the **complete** parameter to load the saved .ldi and .cmd files. If this parameter is omitted, the command will load only the sharable settings in the .ldi file.

The **config load** command discards the current configuration and institutes the loaded configuration; it does not merge the loaded configuration with the pre-existing one. The new configuration settings are then stored in 9.256/CFG/config.ldi.

Keep in mind that the .ldi file includes the unit s password, and if you load the configuration on another unit, you will change its password. If you want to load a traffic configuration on another unit without changing the password, use the class load command instead of the **config load** command.

> **Note:** The PacketShaper image version is stored in the .ldi file if it was set in PolicyCenter. If the image version on a unit is different from the image version stored in an .ldi file you are loading, you may see an image configuration error message after issuing the **config load** command in local mode. You can clear the error by giving the **setup version none** command. The error message does *not* appear in shared mode.

Command Change History

Release	Modification
8.3.1	**[complete]** parameter added, which can load both .cmd and .ldi files.

config reset

For PolicyCenter/ PacketShapers in Shared Configuration Mode

When issued from the command-line interface of an individual PacketShaper, this command disables the unit's connection to the PolicyCenter directory server, returning the unit to local mode and setting the unit's sharable attributes to their factory-default state. The **config reset** command will not remove a unit entry from the PolicyCenter directory server, and the unit's non-sharable settings (IP address, DNS and management port settings, etc.) will not be changed. To completely remove the unit entry from PolicyCenter, use **unit clean**.

When you issue this command from the PolicyCenter command-line interface, communication between PolicyCenter and the directory server will be disabled. With this connection disabled, PolicyCenter will no longer be able to contact PacketShapers in shared mode. To restore the connection between PolicyCenter and the directory server, use **config setup**.

```
config reset
```

> **Note:** If you want to return a unit to local mode *without* clearing the unit's sharable attributes, use **config unset**, instead. You may restore a unit's previous PolicyCenter configuration at any time by resetting its connection to the directory server with the **config setup** command.

config save

Save the current configuration's sharable settings in an .ldi file and its nonsharable settings in a .cmd file.

```
config save [<cfg-path>] <file> [unit]
```

<cfg-path>	To save a PolicyCenter configuration, specify the path of the configuration you want to save.
<file>	Specify a filename. The .ldi and.cmd extensions are automatically added to the configuration file name. This command saves the files to the 9.256/ folder or the PolicyCenter directory.
[unit]	When issuing this command from PolicyCenter, you can include the **unit** parameter to save a unit's local sharable and nonsharable settings. If this parameter is omitted, the **`config save`** command will save a configuration's inherited and local settings.

This command can save the traffic tree, partitions, policies, host lists, basic settings (such as shaping and traffic discovery), security settings (such as passwords and login access protocols), SNTP, and Syslog settings, DNS server, and gateway addresses, domain names, time zones, and network interface settings. Use the setup show command to see a list of sharable and nonsharable settings that are stored in the configuration files.

The **config save** and **config load** commands are useful for experimenting with different configuration settings. For example, you can save your current settings, make changes to the configuration (such as create new partitions or policies), and then return to the original configuration if you prefer it. You can create as many configurations as you like.

This feature can also be used to share configurations with other units. You can download the configuration files to a local workstation, upload them to another PacketShaper unit, and then activate the configuration with the config load command.

> **Note:** Keep in mind that the .ldi file includes the unit s password, and if you load the configuration on another unit, you will change its password. If you want to load a configuration on another unit without changing the password, use the class load command instead of the **config load** command.

Command Change History

Release	Modification
8.3.1	Command modified to create both .ldi and .cmd files.

config secure

Issue this command to enable or disable Secure LDAP communication between PacketShaper and PolicyCenter.

```
config secure on|off
```

config setup

Enables directory server access to PolicyCenter and configures the unit to shared mode.

```
config setup <ldap_host>[:<port>] [secure | unsecure] [<directory_server_password>]
[convert]
```

Where:

<ldap-host>	DNS name or IP address of a PolicyCenter directory server
<port>	TCP port number to connect to on the directory server
secure|nonsecure	Specify **secure** to establish a secure LDAP connection between the PacketShaper and the PolicyCenter directory server, or specify **nonsecure** for a standard LDAP connection.
<directory_server_password>	Password for the PolicyCenter directory server.
[convert]	Specify the **convert** option to convert the unit's existing configuration into a new PolicyCenter configuration with the same attributes and values. Because the unit's new PolicyCenter configuration will be based upon its previous configuration, the unit will continue to operate the same in PolicyCenter as it did in local mode. If you do *not* select the **convert** option, the unit's new PolicyCenter configuration is cleared, and will have default settings only.

If you previously issued the command config unset to disable communication between PolicyCenter and the directory server, you can issue the command `config setup <ldap_host>[<:port>] [secure | unsecure] [<directory_server_password>]` from the *PolicyCenter* configuration (the configuration for the PolicyCenter software) to restore communications between PolicyCenter and the directory server. Note that this use of the **config setup** command doesn't support the *convert* option.

config unset

For PolicyCenter/ PacketShapers in Shared Configuration Mode

This command disables directory server access for a unit, and returns the unit to local mode. The **config unset** command removes a unit entry from the PolicyCenter directory server, so the PacketShaper no longer appears on the PolicyCenter Configurations tab, but allows the unit to retain its last PolicyCenter configuration after it returns to local mode.

```
config unset
```

date

View or set the date and/or time. When initially setting the date and time, use **setup timezone**.

```
date [<yyyymmddhhmm>[<.ss>]]
```

Note that this command has the same functionality as the **setup date** command.

> **Note:** You should always do a system reset immediately after changing the date so that the underlying time-sensitive scheduled operations of the PacketShaper can be correctly initialized.

dns lookup

List the IP address(es) associated with a domain name. PacketShaper keeps the mapping data up to date so that when a site changes an IP address, the matching rule knows about the change.

```
dns lookup <hostname>
```

If the name that you enter is different from the canonical or official name, the canonical name record (CNAME) is displayed at the end of the address list. A canonical name record defines an alias for the official host name, facilitating the transition from an old name to a new name.

Some sites return multiple addresses to a lookup query. The PacketShaper classification process compares the traffic flows to the address lists when looking for a match.

dns names

List all domain names and addresses that are configured in traffic class matching rules.

```
dns names
```

Domain Name	IP Address	TTL	Age	RQSNCRP	Error
(luna-corp.bluecoat.com)...					
	192.168.0.33	3600	647	Q	
(m10-pat-corp.bluecoat.com)...					

```
                192.168.0.207          3600      427   Q
percy.xyz.com   (204.202.49.73)        86400     12512 Q
```

The resolved values are shown in parentheses. The other columns in the output are described below.

TTL: The time interval that the DNS entry may be cached before the source of the information should again be consulted.

Age: The time, in seconds, since PacketShaper received the last name refresh.

R: If a name server cannot be reached, the entry's retry count is incremented. This is a high-level retry, and each one may include multiple queries to each name server. If the retry value is greater than 9, an asterisk is displayed in this column. If the retry value is zero, nothing is displayed in the column.

Q: Displays a Q if PacketShaper sent a query and received a response for the name.

S: Displays an S if PacketShaper learned the name's address (or vice versa) by spying on DNS traffic instead of making a query.

N: The number of successful responses received since the one containing this address. If the value is 0, nothing is displayed in the column.

C: The number of responses received before getting one without any new addresses. This is the length of a round-robin cycle. If the value is 1, nothing is displayed in the column.

R: The number of matching rules that refer to this name. It will be incremented by one while a name is being resolved. If the value is 1, nothing is displayed in this column.

P: Displays a P if PacketShaper is currently resolving this name.

Error: Shows the problem (if any) encountered by the last refresh attempt. Some possible errors are:

name not found: The authoritative server for this domain has no such name.

server offline: The resolver could not reach the authoritative name server, either directly or indirectly through the locally-configured name servers.

rqst refused: The name server knows (or might know) but won't tell you.

no data record: The name exists, but does not have an address (or vice versa).

internal error: The name server is not functioning.

dns refresh

Clear the resolved DNS values that is, names and IP addresses in the names database. The entries then are repopulated at the next ten-second polling interval.

```
dns refresh
```

Immediately after executing **dns refresh**, if you use the dns names command, the resolved values will be listed as <unknown> in the output. These entries are repopulated at the next polling interval.

dns rlookup

Find the host name associated with an IP address.

```
dns rlookup <ipaddress>
```

dns servers

List the DNS servers configured on the system.

```
dns servers
```

dns trace

This is a troubleshooting command that should be used only with the guidance of Customer Support.

email queue

Display or delete messages in the email queue.

```
email queue show|display <message-id>|delete (<message-id>|all)
```

Examples:

To display the email queue:

```
email queue show
```

To display the contents of message 1:

```
email queue display 1
```

To delete all messages in the email queue:

```
email queue delete all
```

Command Change History

Release	Modification
11.5.1	Command introduced

email retry

Deliver mail immediately, rather than waiting for next retry.

```
email retry
```

Command Change History

Release	Modification
11.5.1	Command introduced

email test

Verify the email configuration by sending a test message to individual recipients.

```
email test <recipient> [<recipient>] [<recipient>] [<recipient>]
```

Command Change History

Release	Modification
11.5.1	Command introduced

event delete

Delete an event and all its registrations.

```
event delete <name>
```

Command Change History

Release	Modification
11.5.1	Command introduced

event email

Add or delete an email recipient for event notifications.

```
event email add [<recipient> ... <recipient>]
```

```
 event email delete [<recipient> ... <recipient>]|all
```

Separate recipients with a space. You can add up to four recipient addresses.

To use the command-prompt mode, use:

```
event email add
```

Command Change History

Release	Modification
11.5.1	Command introduced

event log reset

Delete current and archived event log files.

```
event log reset
```

Command Change History

Release	Modification
11.5.1	Command introduced

event log status

Display information about current and archived event log files, such as their location, current capacities, and limitations.

```
event log status
```

Command Change History

Release	Modification
11.5.1	Command introduced

event new

Define a new event. When you define an event, you specify a measurement variable in an expression– that is, the condition for which you want to be notified. In addition, you can define a default event-checking interval. The maximum number of events that can be defined is 32. Defined events are not active until registered.

To initiate the command-prompting mode use:

```
event new
```

You may exit 'event new' at any time by typing 'exit'
Name of the event: **WebQoS**

Type of object to be tested: Link, Partition, or traffic Class: (class):
Measurement Engine variable to be tested: **tcp-conn-aborts%**
Default checking interval [1m,1h] (1m):
Enter a relational operator. When you register this event later,
you will supply a threshold on 'tcp-conn-aborts%' that triggers the event.
The event can be triggered when 'tcp-conn-aborts%'
becomes >, >=, <, or <= the threshold.
Relational operator (>, >=, <, or <=) (>):

As an alternative to the prompting mode, you can use a single command line to create an event, as follows:

```
event new <name> <expression> [<default checking interval>]
```

<name>	Event names must begin with an alphabetic character and contain only alphabetic characters, numbers, and underscores, up to a maximum of 32 characters. Note that you cannot use hyphens in event names.
<expression>	The expression specifies the condition that will be checked and requires adherence to the following syntax: *<variable>*[.*<object type>*] *<relational operator>* *<constant>* Where: *<variable>.<object type>* is one of the PacketShaper measurement variables with an appended object type that is, a link, partition, or class. This object type is required for most variables those that are common to link, partition, and class objects. Some variables are unique to the object type. For example, peak-excess-bps is relevant only to partitions, so it does not need the object-type qualifier in this syntax. Later, when you register the event, you will supply a specific name for the object type. For a list of measurement variables, use the **measure show** command. *<relational operator>* is one of the following: <, <=, =, >=, or >. *<constant>* is a placeholder for the threshold value. Use the $n syntax for example, $1 or $2. When you register an event, you supply a value that is substituted for this constant in the expression. Example of an expression: tcp-rtx-pkts% > 30
[<default checking interval>]	The default frequency that PacketShaper will use to check for this event. When you register this event, you can substitute a different interval. For standard PacketShaper units, you can specify **1m** (one minute) or **1h** (one hour).

Examples:

```
event new NetworkInefficiency tcp-efficiency%.link<$1 1m

 event new WebQos tcp-conn-aborts%.class>$1 1h
```

Command Change History

Release	Modification
11.5.1	Command introduced

event register

Initiate event-checking and notification for an event. The maximum number of events that can be registered at one time is 32. To use the command-prompting mode, simply use event register, otherwise use the following command syntax.

```
event register <event name>(<object>,<threshold>,<re-arm>) [<checking interval>] [email]
[trap] [syslog] [limit=<n>]
```

<event name>	An existing predefined or user-defined event
<object>	The name of a link, partition, or class that is relevant to the event definition
<threshold>	The value used to trigger event notification. The value is substituted in the event's expression, which you defined with the **event new** command. If the condition in the expression occurs, it triggers the event notification that is registered for the event.
<re-arm>	The value that tells PacketShaper that it's okay to once again send event notifications. After the initial notification occurs for the threshold crossing, additional event messages traps, email, or syslog will not be sent until the re-arm condition occurs. The purpose of the re-arm value is to prevent excessive event notification.
[*<checking interval>*]	The frequency at which this condition should be checked. For standard PacketShaper units, you can specify **1m** (one minute) or **1h** (one hour).
[email] [trap] [syslog]	The notification mechanism for this event email, trap, or Syslog.
[limit=*<n>*]	The number of notifications to be sent within the 24-hour period from midnight to midnight. If you omit this option, the number of notifications is limitless.

Example:

```
event new WebQos tcp-conn-aborts%.class>$1 1h

 event register WebQos(inbound/outside/http,70,50) 1m email limit=20
```

Note that in the above example, the event was defined with a default interval of one hour. When the event was registered, the specific class was identified with a threshold of 70%, a re-arm level of 50%, a 1-minute interval, and a limit of 20 notifications within a 24-hour period.

When an event exceeds the predefined threshold value, the event is in violation and the PacketShaper will automatically send out notification. PacketShaper will also send a notification when the re-arm level is crossed, allowing you to be alerted automatically when the event has been cleared.

Command Change History

Release	Modification
11.5.1	Command introduced

event reset

Reset the user events system. This command removes all user-defined events and unregisters all events (user-defined and predefined).

```
event reset
```

> **Note:** Issuing the **event reset** command from the PolicyCenter command line interface can incorrectly trigger an error message stating that the operation failed, even if the operation executed correctly.

Command Change History

Release	Modification
11.5.1	Command introduced

event show

Display email notification recipients, available events (both user-defined and predefined), registered events, and their status.

```
event show
```

Command Change History

Release	Modification
11.5.1	Command introduced

event test-email

Verify the event email configuration with a test email.

```
event test-email
```

Command Change History

Release	Modification
11.5.1	Command introduced

event unregister

Stop checking an event.

```
event unregister <registration-id>|all
```

Use **event show** to display the registration IDs.

Command Change History

Release	Modification
11.5.1	Command introduced

exit

Log out of a PacketShaper connection.

```
exit
```

help

List available commands. Specify a command to view its syntax and usage details.

```
help [<command>]
```

highav add

Define an access router for the access-link monitoring (high availability) feature. This feature allows PacketShaper to deal with imperfect load-balancing and has the ability to respond to the occurrence of WAN link failure. When high availability is enabled, PacketWise can adjust partitions appropriately to prevent overloading any given WAN link and to account for lost available capacity due to router or link failure. High availability has two modes: basic and advanced.

```
highav add <address> <community>
```

where

<address>	IP address of the router
<community>	SNMP community string (password) for the router

Example:

```
highav add 10.10.10.10 pAss4WoRD
```

Command Change History

Release	Modification
11.6.1	Command introduced

highav community

Change the community string of a high availability router. Use this command when the community string changes after you have already defined the router with the highav add command.

```
highav community <address/sysname> <community>
```

where

<address/sysname>	The router s IP address or system name
<community>	New SNMP community string (password) for the router

Command Change History

Release	Modification
11.6.1	Command introduced

highav delete

Remove an existing router from the high availability configuration.

```
highav delete <address/sysname>
```

Command Change History

Release	Modification
11.6.1	Command introduced

highav disable

Disable link monitoring (basic mode) as well as link overload protection (advanced mode, if enabled).

```
highav disable
```

Command Change History

Release	Modification
11.6.1	Command introduced

highav enable advanced

Enable two high availability features: link monitoring/resizing (as in basic mode) and link overload protection. With the link monitoring/resizing feature, the PacketShaper polls the configured router(s) every 30 seconds to assess the status (link up or link down) of the WAN link interfaces. If a link goes down, PacketShaper will automatically adjust the total available capacity by subtracting out the capacity of the down link. With link overload protection, PacketShaper can help prevent the overloading of an interface. PacketShaper will use SNMP polling to access the actual throughput of each configured WAN link interface. If an interface approaches its configured capacity, PacketShaper will pace the traffic sent through that interface to prevent overloading the link and reduce the number of retransmissions. This is accomplished by adjusting the size of the Inbound and Outbound partitions.

```
highav enable advanced
```

To turn off the advanced mode of high availability, use the highav disable command.

Command Change History

Release	Modification
11.6.1	Command introduced

highav enable basic

Enable the link monitoring/resizing high availability feature. When this feature is enabled, PacketShaper polls the configured router(s) every 30 seconds to assess the status (link up or link down) of the WAN link interfaces. If a link goes down, PacketShaper will automatically adjust the total available capacity by subtracting out the capacity of the down link.

```
highav enable basic
```

Suppose you have two routers, A and B. Router A has two 200K interfaces and Router B has one 100K interface. The total available capacity is 500K (unless you have set up an override—see highav override). Now suppose one of Router A's 200K links goes down. With basic high availability enabled, PacketShaper will not only detect the down link, it will also automatically reduce the total available capacity by the capacity of the down link (500K minus 200K = 300K).

To turn off the basic mode of high availability, use highav disable.

Command Change History

Release	Modification
11.6.1	Command introduced

highav interface add

Define the WAN link interface used on a previously-defined access router.

```
highav interface add <address> <interace number/name> <inbound-bps> <outbound-bps>
```

where

<address>	The router's IP address or sysname
<interface number/name>	The name (ifname) or index number (ifindex) that identifies the interface. Examples of interface names are *ethernet 3/1* and *serial 0/1.* It is recommended that you identify the interface by name, not index, because ifnames are unique and persistent while index numbers can change dynamically. If you are using Cisco IOS v12.1 or above and have configured the router to make the ifindex persistent, you can safely identify the interface by index number. Note that ifname was not available in Cisco IOS before v11.1. Instructions for finding the ifName and ifIndex values for Cisco router interfaces
<inbound-bps>	Maximum inbound throughput that is expected to pass through the interface. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions).
<outbound-bps>	Maximum outbound throughput that is expected to pass through the interface

Adding an interface will increase the router's available bandwidth unless you have set override values. The lowest value (override versus sum of interfaces) takes precedence. For example, suppose a router has two 400K interfaces and you have set an override of 600K. If you add another 200K interface, the override will take precedence (in other words, the router's available bandwidth will still be 600K). Make sure that you adjust your override after adding a new interface.

Command Change History

Release	Modification
11.6.1	Command introduced

highav interface delete

Delete a previously-defined interface from the high availability configuration.

```
highav interface delete <address> <interface number/name>
```

where

<address>	The router's IP address or sysname
<interface number/name>	The name or SNMP index number of the interface you want to remove

Deleting an interface may reduce the router's available bandwidth, depending on the override value. For example, suppose a router has two 400K interfaces and you have set an override of 600K. If you then delete an interface, the router's available bandwidth would be reduced to 400K; the override would be ignored since it's greater than the sum of the router's interfaces.

Command Change History

Release	Modification
11.6.1	Command introduced

highav interface modify

Modify the settings for a previously-defined WAN link interface.

```
highav interface modify <address> <interace number/name> <inbound-bps> <outbound-bps>
```

where

<address>	The router's IP address or sysname
<interface number/name>	The name or SNMP index number whose settings you want to modify
<inbound-bps>	Maximum inbound throughput that is expected to pass through the interface. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions).
<outbound-bps>	Maximum outbound throughput that is expected to pass through the interface

Command Change History

Release	Modification
11.6.1	Command introduced

highav override

Configure the inbound and outbound speed of the router. When an override is set, PacketShaper uses this speed for calculating the WAN link capacity for the router, as opposed to using the sum of the interfaces.

```
highav override <address> {<inbound-bps> <outbound-bps>} | none
```

where

<address>	The router's IP address or sysname
<inbound-bps>	Maximum inbound and outbound throughput that is expected to pass through the router. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), G (billions).
<outbound-bps>	To remove the override, use **none**.

|none

This optional approach might be used in a situation with multiple WAN access line interfaces on a router. If you don't expect to get perfect load balancing between the interfaces, you can configure a smaller value for the router than for the sum of the interfaces. If both interfaces are up, PacketShaper would use the override value for the router when calculating the WAN access line capacity available for the router. If one of the interfaces goes down, PacketShaper would use the capacity configured for the active interface (the values configured with the highav interface add command).

Command Change History

Release	Modification
11.6.1	Command introduced

highav show

Show current high availability configuration and status. The output indicates the overall high availability capacity as well as the settings of each interface and router.

```
highav show
```

```
          High Availability: Mode = Basic
                 Access Set: In 500k Out 500k
Total Available Capacity: In 500k Out 500k

          Router Address: 192.168.176.5
                 Active: yes
                SysName: testnetrouter.bluecoat.com
       Override Capacity: No Override Set

              Interface: + ET0(1) speed: 10.0M
     Interface Capacity: In 200k Out 200k

          Router Address: 192.168.176.2
                 Active: yes
                SysName: router1
       Override Capacity: In 300k Out 300k

              Interface: + ET0(1) speed: 10.0M
     Interface Capacity: In 200k Out 200k

              Interface: + ET2(3) speed: 1.5M
     Interface Capacity: In 200k Out 200k
```

The table below describes the output.

High Availability	The current high availability mode (basic, advanced, or disabled)
Access Set	Inbound and Outbound access link speed. In basic mode, these values are the same as Total Available Capacity. With link overload protection (a feature of advanced mode), these values are based on actual throughput observed through SNMP polling. More detailed information about observed minimum values are also listed.
Total Available Capacity	The total bps available based on the values configured for the interfaces and routers. It is the sum of the routers' capacities. A router's capacity is determined by the values set with the highav override command or by summing all the interfaces' capacities (if no override has been set). When high availability is enabled and a link becomes inactive, the Total Available Capacity will reflect this reduction of available bandwidth (that is, the inactive link's capacity will be subtracted out, assuming it is less than the override value).

Router	The router's IP address and sysname that were configured with the highav add command, the router status (active vs. inactive), and the override capacity (if one was set with the highav override command).
Interface	Interface name, SNMP index number, and the inbound and outbound capacities that were configured with the highav interface add command. If you see *Unknown* for the interface name, your router's OS may not support the ifname variable. For example, ifname was not available in Cisco IOS before v11.1. A "+" indicates the interface is active; a "-" indicate the interface is inactive. If advanced mode is enabled, the actual bps throughput (based on SNMP polling) is listed.

Command Change History

Release	Modification
11.6.1	Command introduced

history

The **history** command displays the last 20 commands that were entered into the command line interface; each command is prefixed by a number. Any command on the history list can be executed by using the !*<n>* command, where *<n>* is the number next to the command on the history list.

```
history
```

For example:

```
 history
```

31: setup show
32: help me dump
33: help class rule
34: help setup secure
35: traffic flow -tuO
36: traffic flow -tIPc /inbound/default
37: setup shaping on
38: setup discovery on
39: traffic tree
40: link show
41: class show /inbound/default
42: traffic bandwidth /inbound
43: help class new
44: hostdb show
45: sys info
46: traffic bandwidth
47: cat 9.256/log/bootlog
48: ls 9.258/diag
49: setup shaping off
50: setup discovery off

Typing **!40** would repeat the **links show** command.

hl add

Add entries to an existing host list. When specifying multiple names and/or addresses, separate each with a space.

```
hl add <hostlist> <host> [<host> ...]
```

where *<hostlist>* is an existing host list name, and *<host>* can be specified in any of the following ways:

Type of *<host>*	Example
Host IP address	192.168.1.10
Range of IP addresses Use a dash – with no spaces – between the low and high address in the range.	192.168.1.100-192.168.1.200
Address of the subnet; the CIDR number specifies the number of constant bits in the address range	192.168.10.0/24
Range of subnet addresses; the CIDR number specifies the number of constant bits in the address range Use a dash between the low and high address in the range. Spaces are not allowed before or after the dash or slash characters.	192.168.10.0-192.168.20.0/24
DNS name **Note:** Do not use domain names if you will be using the host list with the host sidedness feature.	www.yourcompany.com

Example:

```
hl add competitors yourcompany.com 192.168.1.00-192.168.1.200
```

hl delete

Remove one or more items from an existing host list.

```
hl delete <hostlist> <host> [<host> ...]
```

where *<hostlist>* is an existing host list name, and *<host>* can be specified in any of the following ways:

Type of *<host>*	Example
Host IP address	192.168.1.10

Type of *<host>*	Example
Range of IP addresses Use a dash – with no spaces – between the low and high address in the range.	192.168.1.100- 192.168.1.200
Address of the subnet; the CIDR number specifies the number of constant bits in the address range	192.168.10.0/24
Range of subnet addresses; the CIDR number specifies the number of constant bits in the address range Use a dash between the low and high address in the range. Spaces are not allowed before or after the dash or slash characters.	192.168.10.0- 192.168.20.0/24
DNS name	www.yourcompany.com

> **Note:** You can only remove hosts the way they were originally added to the host list. For instance, suppose you add a host to the host list by specifying a single IP address. The only way to remove the host is by specifying the single address. You cannot remove this host by entering a range of addresses, a subnet, or a range of subnet addresses.

hl new

Create a host list by defining a unique name and specifying the DNS names, IP addresses, and/or subnets that should be included in the list. You can combine names and addresses in the same list. When specifying multiple names and/or addresses, separate each with a space.

```
hl new <hostlist> [<host> [<host> ...]]
```

where *<hostlist>* is a descriptive name, up to 127 characters; the slash (/) and backslash (\) characters may not be used.

The *<host>* can be specified in any of the following ways:

Type of *<host>*	Example
Host IP address	192.168.1.10
Range of IP addresses Use a dash – with no spaces – between the low and high address in the range.	192.168.1.100- 192.168.1.200
Address of the subnet; the CIDR number specifies the number of constant bits in the address range	192.168.10.0/24

Type of *<host>*	Example
Range of subnet addresses; the CIDR number specifies the number of constant bits in the address range Use a dash between the low and high address in the range. Spaces are not allowed before or after the dash or slash characters.	192.168.10.0-192.168.20.0/24
DNS name **Note:** Do not use domain names if you will be using the host list with the host sidedness feature.	www.yourcompany.com

Host lists are useful when creating classes based on hosts, retrieving host accounting data, and defining exception lists for adaptive response host agents.

The **hl new** command accepts any addresses and/or names that are syntactically correct. It does not validate the existence of the entries.

To add entries to the host list after it's created, use the hl add command.

Examples:

```
hl new BigGifs www.yourcompany.com 192.168.0.116

hl new insidelist
```

hl override

For PolicyCenter / Units in shared mode only

Override an inherited host list by creating a local copy of the list in the current configuration.

```
hl override <list-name>
```

You must make a local copy of an inherited host list before you can change the host list on the child configuration.

hl refresh

Update the host lists with the latest data from the DNS server.

```
hl refresh
```

hl resolve

Display the addresses that are mapped to a particular host list name.

```
hl resolve <hostlist>
```

Example:

```
hl resolve BigGifs
```

ldap:///biggifs,ou=hostlists,ou=m10-pat,ou=pscfg,o=bluecoat.com: 198.3.99.199, 192.168.0.116, 204.71.177.35

hl rm

Remove a host list from the directory configuration.

```
hl rm <hostlist>
```

Host lists cannot be removed if they are currently being used (for example, in a class matching rule or a host side list).

hl show

Display a list of all defined host lists or show the details of a specific host list.

```
hl show [<list_name>]
```

To show all host lists and all host values:

```
hl show *
```

Host values are listed alphabetically or by top-level domain order.

host accounting categories

Create the names of categories to be used in the host accounting feature. For example, an ISP can create categories such as premium, standard, or free, and then assign traffic classes to the appropriate category. This would allow the service provider to charge different rates for premium and standard traffic and discount traffic to free services (such as the ISP support computers).

```
host accounting categories none|<category> [<category> ...]
```

You must create all your categories at once, separating each name with a space. The total string of category names must be 200 characters or less. Forty four is the maximum number of categories you can create, assuming sufficient resources on your PacketShaper.

Note that you cannot issue this command while the measurement engine is in the process of starting or resetting.

After creating categories, use the **class category** command to assign traffic classes to the categories.

You cannot selectively add or delete categories. If you later want to modify your category list, you must use the **host accounting categories** command and specify the complete list of categories you want to have. You will then need to reset the unit to create the new categories in the measurement database. Note that the host accounting measurement data will be cleared when you create categories, so make sure you retrieve your measurement data before adding categories.

Each category name is actually the name of a measurement variable. If you give the **measure show** **host accounting** command, you will see the host accounting categories listed as variables.

To remove all categories, use:

```
host accounting categories none
```

host accounting enable

Enable or disable recording of per-host accounting data in the measurement engine database. The host accounting feature allows you to store and retrieve measurement data per IP address, without having to create a traffic class for each user. This feature is especially useful with dynamic partitions, which create subpartitions for each IP address or subnet. You can use host accounting to see total bytes sent and received per IP address.

```
host accounting enable <mode> [<interval-minutes> [<max-samples>]]
```

<mode>
Specify the host location for which you want to record data:

inside record data for inside hosts

outside record data for outside hosts

both record data for inside and outside hosts

none turn off host accounting

<interval-minutes>
Number of minutes between each recorded sample (the default is 10, the minimum is 1, and the maximum is 1440 minutes)

<max-samples>
Maximum number of samples that can be stored in the host database (default is 1,000,000). This value needs to be greater than the concurrent host limit on your unit (this limit varies by model). *<max-samples>* may require some experimentation. Try a large number (such as 3,000,000) and see if the unit stores host data for a sufficient length of time. If it stores only three weeks of data (and you need it to store a month's worth), you'll need to increase the *<max-samples>* value. Bear in mind that the larger number of samples you store, the more disk space you ll need.

Note that you cannot issue this command while the measurement engine is in the process of starting or resetting.

The building of the host accounting measurement data file can take awhile; the more categories and samples you have, the longer it takes to build the file. While the file is building, you will not be able to issue any commands in the current remote login session. If you open another session while you re waiting, you can issue any command except for the **measure show** command.

Note that any pre-existing host accounting data will be cleared when the data file is built, so make sure you retrieve your measurement data before enabling host accounting (see measure dump).

After you enable host accounting host accounting measurement data will not begin recording again until the next full interval. For example, assume *<interval-minutes>* is 2 and you reset the unit. When you give the **measure show** command, the message indicates "Measurement engine is waiting until 15:03 to start." At 15:03 all other measurement groups will begin recording, but host accounting will not begin recording until 15:04 (the next interval).

To disable recording, use:

```
host accounting enable none
```

host accounting retrieve

Retrieve host accounting data. You can display the data on the screen or save it into a comma separated value (.CSV) file. You can select hosts by IP address, CIDR subnet, or membership in a host list.

```
host accounting retrieve [dns] <ip-addr>|<subnet>/<cidr>|<hostlist>|all from <start-date-time> to <end-date-time> [into <file>|to <file>]
```

[dns]	Specify dns to include DNS names, when available, in the output (instead of IP addresses)
<ip-addr>	Designate the hosts you want to retrieve data for, using one of the following specifications:
<subnet>/*<cidr>*	*<ip-addr>* – host IP address
<hostlist>	*<subnet>*/*<cidr>* – the name of the subnet; the CIDR number specifies the number of constant bits in the address range
all	*<hostlist>* – the name of a host list file
	all – all entries in the time range
from *<start-date-time>*	Specify the starting and ending date and/or time to retrieve data for. The *<start-date-time>* and *<end-date-time>* are required parameters you must specify a date, a time, or the date and time. If a time is omitted, midnight is assumed; if a date is omitted, today's date is assumed. Dates and times can be entered in the following formats:
	M/D (for example, 5/3 midnight of May 3)
	M/D HH:MM (for example, 5/3 13:15 1:15pm on May 9
to *<end-date-time>*	HH:MM (for example, 9:00 9am today
	You can also specify a relative date for example, -7 for 7 days ago.
	Note: If the end date is after today's date, PacketShaper assumes the date you meant was last year's date. If the end date is before the current date and after the start date, PacketShaper will display the requested data (if any exists).

The **into** literal dumps the records to the file named *<file>*. If *<file>* already exists, the records will be appended to the existing file.

into *<file>* / to *<file>*
The **to** literal also dumps the records to the file named *<file>*, but it overwrites the contents of *<file>* if it already exists.

If no path is specified, the file is stored in the current folder (9.256/, by default).

If **into *<file>*** or **to *<file>*** is omitted, the records appear on the screen.

The output contains, for each specified host, a comma-separated-values list of the total bytes recorded over the time period as well as the total for each category.

```
host accounting retrieve dns all from 14:12 to 14:14
```

```
# 26-Jul-2001 14:12:00 to 26-Jul-2001 14:14:00
 host , bytes , web , overhead
 r2.us.rmi.yahoo.com ,9085,9085,0
 ck101.rmi.yahoo.com ,377,377,0
 store.yahoo.com ,1684,1684,0
 ...
 ...
 10.7.6.62 ,42,0,0
 pal.ads.vip.sc5.yahoo.com ,1289,1289,0
```

Or, to see the total usage for the month of March for all of the hosts in a certain subnet:

```
host accounting retrieve 192.168.1.0/24 from 3/1 00:00 to 3/31 23:59
```

> **Note:** You can also use the measure dump command to retrieve host accounting data.

host accounting show

List the recording parameters for host accounting that were set with the host accounting enable command. It also lists all the categories that were created with the host accounting categories command, along with the traffic classes assigned to each category (with the class category command).

```
host accounting show
```

```
                Recording Mode: both
                 Sample Interval: 5 minutes
              Allocated Samples: 1000000

 Category      Traffic Class
 ---------------------------------------------
 web
```

/Inbound/http
/Outbound/http
overhead
/Inbound/ICMP
/Outbound/ICMP

hostdb cache

Display the current connections in the IP classification-accelerator cache, or list the class name and matching rule of a specific IP address in the cache. The cacheing feature stores qualified IP address-based classes in a cache, thereby increasing the speed in which PacketShaper classifies flows on the inside of the unit.

```
hostdb cache [<ipaddress>]
```

If you don t specify an IP address, the output lists the current connections in the host cache:

```
hostdb cache
```

IP Address	Direction	Class
10.7.38.100	inbound	mysite.org
10.7.38.100	outbound	mysite.org
10.7.6.81	outbound	default
10.7.39.12	outbound	default
10.7.6.12	outbound	default
10.7.40.1	outbound	default

The following example lists details about a specific IP address (10.7.38.100) in the cache. It indicates the class name and match rule that has been cached for a particular host.

```
hostdb cache 10.7.38.100
```

Traffic Class: /Inbound/10.7.38.0/CUSTOMER/mysite.org
Match rule used for this host, 10.7.38.100:

[1] inside host 10.7.38.100 any port IP
outside any host any port
Traffic Class: /Outbound/10.7.38.0/CUSTOMER/mysite.org
Match rule used for this host, 10.7.38.100:
[1] inside host 10.7.38.100 any port IP
outside any host any port

hostdb info

Display the host IPv4 or IPv6 address, user name, average and current connections, current guaranteed and excess bandwidth, and throughput information. **Hostdb info** shows more detailed rate information than the **hostdb show** command.

The host database is a record of all hosts that have active connections through the unit. Once a host closes its connection, the host will be purged from the database. In addition, the unit will clear host entries if they aren't active for approximately ten minutes. Thus, the hostdb is a real-time list of hosts.

`hostdb info [<sort-switch>] [<number-switch>] [<switch>|all] [<host_addr>|<host_name> [<mask>]]`

[*<sort-switch>*]	Specify one of the following switches: **-sf** or **sortfpm** (sort hosts by flows per minute in descending order) **-sp** or **sortfail** (sort hosts by rate of failed new TCP connections) **-sr** or **sortrate** (sort hosts by current rate in descending order) **Note:** Because the host database changes even as the **hostdb** info command is executed, the returned list of hosts will not always display in decending order.
[*<number-switch>*]	Specify the number of hosts to display: -n *<number>*
[*<switch>*|all]	Specify one of the following switches: **-a** or **active** (currently connected hosts only) **-o** or **outside** (hosts outside the unit) **-i** or **inside** (hosts inside the unit) **-u** or **unknown** **Note:** The **active** switch can be used in conjunction with the other switches. Or, use **all** to show information for all hosts in the database. When you omit the ***<switch>*|all** parameter, the command displays hosts that have accessed the unit within the last five minutes, but may not be currently connected.
[*<user-switch>*]	Specify a user name to display hosts for: -user *<user_name>* If the user name contains spaces, it must be enclosed in quotes. Specifying the domain name is optional.
[*<host_addr>*|*<host_name>*]	Specify either a host's IP address or name.
[*<mask>*]	The host's subnet mask

The following example shows information for all the hosts in the database:

```
hostdb info

IP Address                    Conn    RTT      Cur    1 Min    Peak --- New Flows Per Minute --- User
                                     to PS     rate     avg    rate    Client  Server  Failed
-------------------------------------------------------------------------------------------------
11.1.2.121                  I  53     908ms     187k    113k    221k     439       0       0     cal\jsmith
11.1.2.122                  I  52     912ms     122k    108k    212k     422       0       0     cal\mjohnson
11.1.2.123                  I  54     910ms     133k    106k    216k     416       0       0     cal\speters
11.1.2.124                  I  51     907ms     169k    110k    218k     439       0       0     cal\bturnbull
11.1.2.125                  I  51     911ms     128k    109k    214k     428       0       0     cal\panderson
11.1.2.126                  I  54     907ms     178k    114k    218k     439       0       0     cal\rcilk
11.1.2.127                  I  52     907ms     169k    110k    218k     439       0       0     cal\dgiordano
11.1.2.128                  I  54     907ms     184k    114k    218k     439       0       0     cal\jgrell
11.1.2.129                  I  50     908ms     179k    114k    219k     439       0       0
cal\egiovanetti
```

The displayed information includes:

IP Address	Identifies the IPv4 or IPv6 host that is connected through the unit. This field is followed by either an I or O, indicating the location of the host (Inside or Outside) relative to the unit.
Conn	The number of connections
RTT to PS	The round-trip time from the host to the PacketShaper
Cur rate	The current rate for the host in Kbps
1 Min avg	A one-minute moving average for the host's rate in Kbps
Peak rate	The highest rate the host's connection has reached. This is the sum of the inbound and outbound traffic, relative to the host.
New Flows Per Minute	Shows the rate of initiation of new flows from this host (as Client) and to this host (as Server). This rate can be limited via the **policy flowlimit** command. The Failed column shows the rate of new TCP connections per minute that the host initiated but failed, either because the host received an immediate RST response or received no response at all. IP addresses with many failed connections are good candidates for more scrutiny; they may be overloaded servers, clients initiating port scans or systems involved as an initiator or recipient of attacks. A "+" next to a host's New Flows Per Minute value indicates load shedding is occurring or has recently occurred. For example, if the Failed column for a host lists 105+, that host has flows that are being shed. For more information about load shedding, see setup loadshedding.
User	User name associated with the IP address. Note that the user awareness feature requires that BCAAA be installed and configured.

> **Note:** Since most web browsers open multiple simultaneous connections, a web policy set to 100 Kbps may actually allow, for example, 400 Kbps *per PC* if the browser is configured to allow four simultaneous connections. This impacts the peak flow numbers for a class.

The following example shows details about a specific host:

```
hostdb info 65.174.190.201

IP Address                  Conn   RTT     Cur   1 Min   Peak --- New Flows Per Minute ---User
                                   to PS   rate    avg   rate    Client  Server  Failed
-------------------------------------------------------------------------------------------
65.174.190.201           O   2     34ms    296k   301k   913k       0       0       0

1 entries matching 65.174.190.201 255.255.255.255
```

In the output of the **hostdb info** command, some of the fields are not populated and list their values as 0. To see information about the recent throughput of known hosts use the hostdb show command instead.

To display the top 5 bandwidth users:

```
hostdb info -sr -n 5


IP Address                  Conn   RTT     Cur   1 Min   Peak --- New Flows Per Minute ---User
                                   to PS   rate    avg   rate    Client  Server  Failed
-------------------------------------------------------------------------------------------
65.174.190.201           O   2     30ms    298k   296k   773k       0       0       0
192.168.0.7              I   8     ---     297k   299k   773k       3       0       0
2001:db8:1234:5678::1    O   0     88ms   1326   3811    140k       0       0       0
192.168.0.175            I   1     ---    1047    779     51k       4      12       0
12.104.153.33            O   0     ---      11     45     63k       0       0       0

5 entries
```

To find an infected host on the network, you can display the top 10 hosts that have the most failed flows during the last minute:

```
hostdb info -sp -n 10
```

To find a host that might be propagating a virus or worm, you can display the top 10 hosts with the most flows:

```
hostdb info -sf -n 10
```

Command Change History

Release	Modification
9.2.1	*<user-switch>* added; User column added to output

hostdb rtostats

Display a list of hosts that have sent premature retransmission timeout (RTO) segments.

```
hostdb rtostats
```

On some server systems the TCP stack uses a short retransmission time interval, which causes premature RTO. This results in unnecessary packet retransmissions on low-speed links, which waste bandwidth.

PacketShaper controls retransmission timeouts for inbound and outbound retransmissions by discarding premature RTO segments. If PacketShaper detects a premature retransmission (retx), it delays that packet for an additional time period, based on measured host latency. If an acknowledgment is received before the retransmission is to be forwarded, the retransmission is discarded.

The **hostdb rtostats** command shows which hosts are experiencing premature RTO. The output also indicates how often PacketShaper "clamped down" on RTO segments.

With this RTO clamping feature, needless retransmissions (shown as *outClampedSegs* and *inClampedSegs* in the **hostdb rtostats** output) are discarded. The term *segment* refers to a TCP datagram. RTO clamping works only when shaping is turned on.

hostdb show

Display the host IPv4 or IPv6 address, estimated access speed, number of speed changes, the number of TCP and UDP flows that a specified host has processed, the amount of time the host has been idle, and the status of the match rule cache.

```
hostdb show [<switch>|all][<host_addr>|<host_name> [<mask>]]
```

[*<switch>*|all]	Specify one of the following switches: **-a** or **active** (currently connected hosts only) **-o** or **outside** (hosts outside the unit) **-i** or **inside** (hosts inside the unit) **-u** or **unknown** **Note:** The **active** switch can be used in conjunction with the other switches. Or, use **all** to show information for all hosts in the database. When you omit the ***<switch>*|all** parameter, the command displays hosts that have accessed the unit within the last five minutes, but may not be currently connected.
[*<host_addr>* / *<host_name>*]	Specify either a host's IP address or name.
[*<mask>*]	The host's subnet mask

```
hostdb show

IP Address                  Side  Speed/Effective TCP/UDP  Time   I O   I R S NC
-------------------------------------------------------------------------------
1.2.3.4                      OUT      0/0           0/0    288s   ? ?   n n n n
1.255.255.255                N/A      0/0           0/0    288s   ? ?   n n n n
10.1.1.16                    OUT      0/19          0/0     23s   ? ?   n n n n
10.1.1.20                    OUT      0/25153       1/0     83s   ? ?   n n n n
10.1.1.27                    OUT      0/510.7k      1/0     20s   ? ?   n n n n
10.1.1.45                    OUT      0/876         2/0    234s   ? ?   n n n n
```

```
10.1.1.120                  out      0/1           0/1      2s   ? ?   n n n n
10.100.99.32                OUT      0/170.3k      1/0     31s   ? ?   n n n n
172.21.0.20                 OUT      0/394.2k      0/0     11s   ? ?   n n n n
172.21.0.84                 out      0/21          0/0    202s   ? ?   n n n n
172.21.1.26                 OUT      0/784         0/0    154s   ? ?   n n n n
2001:db8:1234:5678::2        IN   1.5M/1.5M        6/1      0s   ? ?   n n n n
2001:db8:1234:5678::1       in      0/0           0/0     30s   ? ?   n n n n
```

The displayed fields include:

IP Address	IPv4 and IPv6 addresses of hosts that have communicated through the PacketShaper
Side	Inside or Outside host. If N/A is displayed, the traffic seen from the host is broadcast or multicast. **Note:** The capitalization (*IN* versus *in* and *OUT* versus *out*) indicates PacketShaper's level of confidence for the host's side. If a packet's source host is seen on the inside interface, there is high confidence the host is on the inside of the PacketShaper (*IN* – all caps). It is then assumed with low confidence that the packet's destination host is on the outside interface (*out* – lowercase). If the host is set to the wrong side, you can override this setting with the hostdb side set command.
Speed	Estimated access speed of the connection in bits per second
Effective	Effective rate
TCP/UDP	Number of TCP connections or UDP sessions currently used by this host. If these fields contain zeroes, the host is not currently communicating through the unit.
Idle Time	Amount of time, in seconds, since a packet was received for/from the host
Cache	The status of the match rule cache: I – entry has been cached for the Inbound direction O – entry has been cached for the Outbound direction ? – currently unknown whether the entry for the corresponding direction is in the cache N – entry cannot be cached for the corresponding direction

The following example shows details about a specific PacketShaper:

```
hostdb show 172.21.0.85
```

```
IP Address: 172.21.0.85 OUTSIDE
 Time since last touched: 8194 secs
 Current References: TCP 0 UDP 0
 Speed:      1.5M Bps
 Effective:  11.4M Bps
```

hostdb side auto

Enable automatic side detection. In this mode (the default), PacketShaper automatically determines whether a host is inside or outside, relative to the PacketShaper. This is appropriate for many network topologies; however, for complex topologies, you can manually override the automatic host side detection and force the placement of certain hosts or subnets on the appropriate side (see hostdb side manual).

```
hostdb side auto
```

After you turn on automatic side detection, the current settings display. For example:

```
Mode: Automatic
 Inside:
 Host list: none

Outside:
 Host list: none
```

hostdb side default

Set host side detection to its default mode (auto mode). If the unit is subscribed to PolicyCenter, the **default** option tells PolicyCenter to delete the setting in the local configuration and inherit from the parent configuration.

```
hostdb side default
```

hostdb side manual

Enable manual side mode; this is necessary when PacketShaper isn't able to automatically detect the correct side for certain hosts. When this mode is enabled, you can force the placement of certain hosts on the appropriate side (inside or outside). This is done by using the hostdb side set command to assign hosts or host lists to the inside or outside. The side lists are not actually used until manual side mode is enabled.

```
hostdb side manual
```

After you turn on manual mode, the current settings display:

```
Mode: Manual
Inside:
 Host list: none
 192.21.18.175
 192.21.18.177
 192.21.18.178-192.21.18.180

Outside:
 Host list: none
```

For any host that isn't assigned to a specific side when manual side mode is enabled, PacketShaper will use its normal mechanism for determining and setting a side. In other words, the sides of all other hosts are detected automatically.

hostdb side reset

Clear the side settings for a particular host or all hosts. The next time PacketShaper sees a flow from that address it will again try to figure out whether the host is inside or outside. This might be necessary if a particular host is seen on the wrong side – you can add the host to the proper side list (inside or outside) and then reset the host so that PacketShaper will rediscover the host and place it on the correct side. To see which side PacketShaper considers a host to be on, use the hostdb show command.

```
hostdb side reset all|<ip-addr>
```

Example:

To clear the side setting for 224.0.1.3:

```
hostdb side reset 224.0.1.3
```

If you immediately issue the **hostdb show** command, you'll see that the entry in the Side column for this particular host has been cleared.

IP Address	Side	Speed/Effective	TCP/UDP	Idle Time	Cache I O	Compress I R S NC
10.10.1.52	in	10.0M/4.5M	3/0	50s	? ?	n y n n
10.1.1.81	out	0/44756	1/0	0s	I ?	y n n n
224.0.1.3		0/0	0/0	38s	N O	n n n n
10.1.5.1	in	0/0	0/0	3s	? O	n n n n
10.7.12.20	out	0/0	0/0	24s	I ?	n n n n
10.7.10.10	out	0/0	0/0	0s	? ?	n n y n
6 entries						

hostdb side rm

Remove a host from the manually configured side list. Use this command if you no longer want a particular host assigned to a side. After you remove a host, PacketShaper will determine and assign a side to the host, using its normal mechanism.

```
hostdb side rm list:<hostlist>|<ip-addr>|<subnet>/<cidr>|hosts|all
```

list: *<hostlist>*	Name of the host list file to be removed from the side list **Note:** This does not delete the list – use the hl rm command if you want to delete the list.
<ip-addr>	Host IP address or a range of IP addresses to be removed from the side list To specify a range, use a dash – with no spaces – between the low and high address in the range (for example, *192.168.1.100-192.168.1.200*).
<subnet> /*<cidr>*	The address of the subnet or a range of subnet addresses to be removed from the side list; the CIDR number specifies the number of constant bits in the address range To specify a subnet range, use a dash between the low and high address in the range (for example, *192.168.10.0-192.168.20.0/24*). Spaces are not allowed before or after the dash or slash characters.
hosts	Removes all individually defined IP addresses, ranges, and subnets (but not host lists)
all	Removes all hosts from the inside and outside lists, including host lists

> **Note:** To verify the host was removed, use the hostdb side show command.

Examples:

```
hostdb side rm 172.17.72.0-172.17.75.0/22

hostdb side rm list:outside_list
```

hostdb side set

Assign hosts to the inside or outside. This is necessary when PacketShaper isn't able to automatically detect the correct side for certain hosts.

```
hostdb side set inside|outside list:<hostlist>|<ip-addr>|<subnet>/<cidr>
```

Designate inside or outside hosts, using one of the following specifications:

list: *<hostlist>*	The name of a host list created with the hl new command; only one host list per side is allowed. **Note:** Host lists are the recommended method of specifying hosts. However, you cannot use a host list that contains domain names.
<ip-addr>	Host IP address or a range of IP addresses To specify a range, use a dash – with no spaces – between the low and high address in the range (for example, *192.168.1.100-192.168.1.200*).

<subnet> /*<cidr>*	The address of the subnet or a range of subnet addresses; the CIDR number specifies the number of constant bits in the address range To specify a subnet range, use a dash between the low and high address in the range (for example, *192.168.10.0-192.168.20.0/24*). Spaces are not allowed before or after the dash or slash characters.

Additional Information

- After you assign hosts to sides, you will need to enable manual side mode with the hostdb side manual command.
- For any host that isn't assigned to a specific side when manual side mode is enabled, PacketShaper will use its normal mechanism for determining and assigning a side.
- To remove a host after you have assigned it to a side, use the hostdb side rm command.
- To view a list of hosts assigned to each side, use the hostdb side show command.
- A maximum of 32 entries can be assigned to the inside and outside. An entry can be a single IP address, a range of IP addresses, a subnet, a subnet range, or a host list. Only one host list can be assigned to a side.

Examples:

In this example, host lists named *inside_list* and *outside_list* were created with the **hl new** command. Inside_list contains a list of hosts and subnets that are known to be on the inside of PacketShaper and outside_list contains the hosts known to be on the outside of PacketShaper. To assign each of the hosts in inside_list an inside designation, use this command:

```
hostdb side set inside list:inside_list
```

To assign each of the hosts in outside_list an outside designation, use this command:

```
hostdb side set outside list:outside_list
```

And then enable manual mode:

```
hostdb side manual
```

Each time you assign IP addresses or subnets with the **hostdb side set** command, the specified hosts are added to the appropriate side – you do not overwrite previous settings. For example, the following two commands will assign two hosts to the outside:

```
hostdb side set outside 192.21.18.172
hostdb side set outside 192.15.17.45
```

However, this rule does not apply to host lists since only one host list is allowed per side. If you assign a host list to a side that already has a host list defined, this list will override the one that was previously defined.

hostdb side show

Display host side settings. The current mode (auto vs. manual) is displayed, along with the list of hosts assigned to each side.

```
hostdb side show
```

Mode: Manual
Inside:
Host list: inside_list
192.21.18.175
192.21.18.177
192.21.18.178-192.21.18.180>

Outside:
Host list: outside_list

hostdb topusers

Determine which hosts or users are consuming the most bandwidth. You can configure PacketShaper to track the Top Talkers (hosts which initiate the most traffic) and Top Listeners (hosts which receive the most traffic).

To display statistics for the top 20 bandwidth users per traffic class either receivers or senders use the following commands:

```
hostdb topusers start <tclass> [talk|listen]

 hostdb topusers stop <tclass> [talk|listen]

 hostdb topusers reset <tclass> [talk|listen]

 hostdb topusers show [<tclass>] [talk|listen]
```

start	Starts tracking top hosts (talkers or listeners) for a traffic class
stop	Stops tracking top hosts for a traffic class
reset	Clears the list of top hosts and restarts the host-tracking process
show	Displays the hosts or users that have used the highest percentage of bandwidth in the class since tracking was started. The list is cleared with the **hostdb topusers reset *<tclass>*** command or when you reset the unit. A host stays on the top-20 list until another host uses more bandwidth, at which point the host may drop off the list entirely or move further down the list. For example, suppose top talkers is turned on for the Inbound/HTTP class, and cnn.com is the top consumer with 22%. If another host, yahoo.com, later consumes more bandwidth than cnn.com, yahoo.com might go to the top of the list and cnn.com would drop lower on the list.

Additional Information

- A total of 32 Top Talkers and Top Listeners (combined) can be enabled at one time. Thus, you can enable both Top Talkers and Top Listeners for up to 16 different classes. Or, if you enable one or the other (just Top Talkers or just Top Listeners), you can track top hosts on 32 different classes.
- For non-IP traffic, PacketShaper does not track sessions or hosts. Therefore, traffic for non-IP protocols (IPX, AppleTalk, NetBEUI, DECnet, FNA, and SNA) will not appear in the Top Talker or Top Listener lists.

- The Group Name column in the output shows one of the user group names to which the user belongs. If a top user belongs to more than one group, an ellipses will appear. To see all the group names for a user, use the setup bcaaa server-test command. Note that the only group names that display are ones for which a user group class has been created. The user may belong to other Active Directory groups, but they will not be listed for the user unless a class exists for that group.

Examples:

To start top talker tracking on the Inbound/HTTP class:

```
hostdb topusers start inbound/http talk
```

To see a list of top talkers in the Inbound/HTTP class:

```
# hostdb topusers show inbound/http talk
Top talker analysis for inbound class HTTP.
Duration: 02:09:18
7 active entries.
User Name           Group Name       DNS Name                                                Percent
IP Address
-------------------------------------------------------------------------------------------------------
----------------------------
N/A                 N/A              a184-84-222-35.deploy.akamaitechnologies.com             52
184.84.222.35
N/A                 N/A              a184-84-222-107.deploy.akamaitechnologies.com            24
184.84.222.107
john.smith          group-sales      No such name                                             7
10.200.10.129
N/A                 N/A              20052lpweb01.redcrossblood.org                           4
174.120.176.2
N/A                 N/A              75.126.14.205-static.reverse.softlayer.com               4
75.126.14.205
N/A                 N/A              server-205-251-203-169.lax3.r.cloudfront.net             2
205.251.203.169
N/A                 N/A              a184-84-222-115.deploy.akamaitechnologies.com            1
184.84.222.115
```

```
# hostdb topusers show

2 active top user sessions.

Direction   Class                T/L       Duration
-------------------------------------------------
inbound     HTTP                 talker    02:13:13
outbound    HTTP                 listener  02:13:13
```

ipfilter

Creates an IP filter that configures a PacketShaper to filter traffic based on IP address.

```
ipfilter discard|onlyaccept|passthrough <device> src|dst <ipaddress>[:<port>] [<mask>]
[prot:tcp|udp|icmp]
```

discard
Configures the IP filter to discard packets arriving on <*device*> with the **src**|**dst** IPv4 or IPv6 address as <*ipaddress*>.

onlyaccept
Configures the IP filter to only accept packets arriving on <*device*> with the **src**|**dst** IPv4 or IPv6 address as <*ipaddress*>.

passthrough
Configures the IP filter to pass through packets arriving on <*device*> with the **src**|**dst** IPv4 or IPv6 address as <*ipaddress*>.

<*device*>
The PacketShaper interface on which the IP filter will act. Depending on which NICs are installed, available interfaces may include the following:

Dual-port NICs:

- slot1_in
- slot1_out
- slot2_in
- slot2_out
- slot3_in
- slot3_out
- slot4_in
- slot4_out

Quad-port NICs:

- slot3_in1
- slot3_out1
- slot3_in2
- slot3_out2
- slot4_in1
- slot4_out1
- slot4_in2
- slot4_out2

To see a list of available interfaces on your PacketShaper, use the **help** command. For example:

help ipfilter discard

src|dst *<ipaddress>* [:*<port>*] [<*mask*>] **[prot:tcp | udp | icmp]**	Specifies whether the IP filter applies to traffic originating from (**src**) or directed to (**dst**) the specified host. <*ipaddress*>[:*<port>*] The IPv4 or IPv6 address of the host, with an optional port number. To specify a wildcard, enter 0.0.0.0 as the IP address and specify a port number or a protocol. Enclose IPv6 addresses in square brackets if a port number is specified (e.g., `[3ffe:b00:1:2::5efe:c000:0201]:8080`). [<*mask*>] The net mask of the specified host **[prot:tcp | udp | icmp]** The layer 4 protocol to filter traffic by To filter all traffic to and from a specified host, you must create an IP filter for the host as both the source (**src**) and destination (**dst**) of network traffic.

You can create up to 2,000 IP filters on a PacketShaper; filter entries are saved in the PacketShaper configuration file.

IPv6 Requirements and Limitations

- The IPv6 address must be a unicast IPv6 address; multicast addresses are not supported.
- The IPv6 address can be either link local or global scope.
- IPv6 prefixes are not supported.
- The IPv6 wildcard address (::) is supported.

Examples

This example creates an IP filter that will discard all traffic on the inside interface of the slot 2 NIC that originates from the host with the IP address 10.1.1.14:

```
ipfilter discard slot2_in src 10.1.1.14
```

This example creates an IP filter that will only accept traffic on the slot 1 outside interface that originates from the host with the IP address 10.10.1.1. All other traffic not otherwise managed by an IP filter is discarded.

```
ipfilter onlyaccept slot1_out src 10.10.1.1
```

This example creates an IP filter that allows all traffic destined for the IP address 10.1.10.1 to pass through the slot 1 inside interface.

```
ipfilter passthrough slot1_in dst 10.1.10.1
```

This example creates an IP filter that allows all traffic that originates from the host with the IPv6 address fe80::200:f8ff:fe21:67cf to pass through the slot 1 inside interface.

```
ipfilter passthrough slot1_in src 3ffe:b00:1:2::5efe:c000:0201
```

This examples creates an IP filter that will allow all traffic destined for port 80 on the inside interface of the slot 2 NIC:

```
ipfilter passthrough slot2_in dst 0.0.0.0:80
```

> **Note:** For each IP filter, PacketShaper assigns a unique eight-digit alphanumeric identifier, such as “DC73DA16”. These identifiers are used to specify an IP filter for removal when using the **ipfilter clear** command. To see a list of all configured IP filters and their identifiers, use the **ipfilter show** command.

Command Change History

Release	Modification
11.10.1	Support added for IPv6
11.5.2	Parameters added to filter by port and layer 4 protocol; ability to specify wildcards
11.1.1	Updated to support PS S-series device names
8.2.0	Command introduced

ipfilter clear

Removes all IP filers (default) or the ip filter that you specify.

```
ipfilter clear [<id>]
```

ipfilter clear Removes all IP filters.

[<*id*>] Removes the IP filter identified by [<*id*>]. To delete a single IP filter from a PolicyCenter sharable configuration, create a draft version of that configuration and then issue the command **ipfilter clear *<id>*** from the draft. The deleted filter will be permanently deleted when you commit the draft.

To see a list of all configured IP filters and their identifiers, use the ipfilter show command.

Examples

This example removes all configured IP filters from the PacketShaper.

```
ipfilter clear
```

This example removes only the IP filter with the identifier “DC73DA16”.

```
ipfilter clear DC73DA16
```

Command Change History

Release	Modification
8.2.0	Command introduced

ipfilter iponly

Configures a PacketShaper to relay only IP traffic.

```
ipfilter iponly on|off
```

on Creates an IP filter that relays only IP traffic (applies to all interfaces).
off (Default). Configures the PacketShaper to relay both IP and non-IP traffic.

Command Change History

Release	Modification
8.2.0	Command introduced

ipfilter show

Shows all configured IP filters.

```
ipfilter show
```

Returned data include the status of the ipfilter iponly setting (whether or not the PacketShaper will relay only IP or both IP and non-IP traffic), the unique identifier of the IP filter, number of hits (instances where the IP filter rule matched a packet), and the configuration of the IP filter.

Example

In this example, the PacketShaper is configured to relay both IP and non-IP traffic, and one passthrough IP filter has been configured on the outside interface:

```
 ipfilter show
 Relay all traffic.
 Exclude Filters: total 1
  [DC73DA16] hits 0 Outside
       src 172.21.1.44 (ffffffff)  --> passthru
 Include filters: total 0

ipfilters MIB:
[ 0] outs                       0 [ 1] onlyAccepts          0
[ 2] onlyExcludes               0 [ 3] nonIpDiscarded       0
[ 4] atalkDiscarded             0 [ 5] ipxDiscarded         0
[ 6] netbiosDiscarded           0 [ 7] snaDiscarded         0
[ 8] fnaDiscarded               0 [ 9] decDiscarded         0
```

Command Change History

Release	Modification
8.2.0	Command introduced

license birthcert

Reinstalls the hardware birth certificate. This operation may be necessary if the birth certificate got corrupted or accidentally deleted. This command requires that the PacketShaper have Internet access.

```
license birthcert
```

Command Change History

Release	Modification
11.2.1	Command introduced

license load

License keys, as well as licenses for add-on components, can be downloaded from the Symantec Licensing Portal, uploaded to the PacketShaper or a web server, and then activated with the **license load** command.

```
license load file://localhost/<file path> | http://<url>
```

where *<file path>* is the name and location of the license key file on the PacketShaper and *<url>* is the URL of the file on a web server. Filenames are case sensitive so make sure to use the proper upper/lower case.

Examples:

```
license load file://localhost/9.256/mylicense.bcl
```

```
license load http://mywebserver.com/licenses/mylicense.bcl
```

Additional Information

- After loading the license key, you will need to reset the PacketShaper to activate the license.
- Use the version verbose command to verify that the new key has been loaded.
- When loading a license key that changes the number of supported traffic classes, you will need to reset measurement data. See measure reset.
- The license activation and expiration dates appear next to **License Status** on the Info tab in the Advanced and Sky UIs.
- If a license will be expiring within the next 90 days, a message (such as *PacketShaper license will expire in 15 days*) appears in the CLI banner and on the Info tab in the Advanced and Sky UIs.

Command Change History

Release	Modification
11.2.1	Ability to load license from a file on PacketShaper or a file on a web server
11.1.1	Command introduced

license update

Connects to the Symantec licensing server and downloads the latest license keys. Use this command to immediately retrieve the license after initially configuring the PacketShaper or after a licensing renewal/purchase. This command requires that the PacketShaper have Internet access.

```
license update
```

Additional Information

- After updating the license keys, you will need to reset the PacketShaper to activate the license.
- When loading a license key that changes the number of supported traffic classes, you will need to reset measurement data. See measure reset.
- The license activation and expiration dates appear next to **License Status** on the Info tab in the Advanced and Sky UIs.
- If a license will be expiring within the next 90 days, a message (such as *PacketShaper license will expire in 15 days*) appears in the CLI banner and on the Info tab in the Advanced and Sky UIs.

Command Change History

Release	Modification
11.2.1	Command introduced

links show

Display the programmed link speeds with current link statistics (current rate, one-minute average, and peak rate).

```
links show

               Interface            Speed     Cur    1 Min    Peak
                                    rate      avg    rate
-------------------------------------------------------
Slot2_in            10Gbps        0        0     556M
Slot2_out           10Gbps        0        0    21.3M
Slot1_in            10Gbps        0        0       0
Slot1_out           10Gbps        0        0       0
```

```
Mgmt                 1Gbps     4154     1477      9.0M
```

If traffic shaping is enabled, the output will also show statistics for each direction. For instance:

```
                Direction             Speed      Cur    1 Min     Peak
                                     rate      avg     rate
-------------------------------------------------------
Inbound               1.0G        0        0     556M
Outbound              1.0G        0        0    21.3M
```

Additional Information

- The Inside and Outside statistics measure the traffic that *enters* the unit through these ports. Localhost traffic is not included in the Inside and Outside statistics.
- The bps values for the management port represent Localhost traffic only.
- Use the setup link command to configure the inbound and outbound link speeds.

ls

Display file listings of a specific folder.

`ls [<directory> | <file>]...`

measure cumvar

> **Note:** The **measure cumvar** command is intended as a troubleshooting tool to be used with the guidance of Customer Support.

Use this command to display a base measurement value. A base measurement variable is the accumulated value that is sampled to generate the measurement engine time-series data.

`measure cumvar <type> <element> <var>`

<type>	The element type of the element whose variable is to be viewed: link, partition, or class
<element>	The name of the element whose variable is to be viewed: link (inbound or outbound), partition name, or class name. This command is valid for "leaf" classes only that is, classes that do not have children. If you specify a non-leaf class, it will display a variable value of zero.
<var>	The variable to be viewed. This may be a simple variable, the high 32 bits of a 64-bit variable, indicated by the suffix "Hi", or a particular histogram value, indicated by the suffix "[*<name>*]", where *<name>* is the histogram index name.

For example, use the **measure show** command to get a list of variables for a measurement type, then use **measure cumvar** to display the value:

```
measure cumvar class outbound/outside/http class-hits
```

Base value of "class-hits" on "class outbound/outside/http" is 18.

measure dump

Display a list of comma-separated measurement values at the command line or redirect the measurement data to a file. When you redirect the data to a file, you can generate graphs using a spreadsheet application.

This command provides four selections for ordering the output (one is required):

- by element
- by event
- by variable
- by time

measure dump by element

List measurement data totals for a specific element, such as a particular link (Inbound or Outbound), partition, class, or host IP address. You can list data for one or more measurement variables for a designated duration (such as 1 day or 2 weeks). The measurement data can be displayed onscreen or stored in a delimited text file. The **by element** option dumps the same data as the **by var** option, except the output is structured differently. (examples that compare **by var** and **by element** output)

```
measure dump [partner <partner_ip>] link|partition|class|host [group day|month|accounting]
[immediate|leaf|all] [<link-name>|<partition-name>|<class-name>|<host-ip-addr>...]
[endtime [<date>]<time>] by element <duration> [to <file>] [sort] <var>...
```

partner *<partner_ip>*	Dumps data for the specified standby partner. If a partner is not specified, data for the current PacketShaper is dumped.
link|partition|class|host	Specify the type of data to be dumped.
[group day|month]	Dumps data from a specific measurement group day, month. If you don't specify a group, PacketShaper will choose the group that appears to most appropriate for the interval and/or duration you have specified. For example, if you specify an interval of 15 minutes, PacketShaper assumes you want to dump the data from the day group. Or if you specify an interval of 2 hours, PacketShaper assumes you want to dump the data from the month group. You can use the **group** option to override the default behavior.

[immediate|leaf|all]	**immediate** – dumps data for the specified class and direct child classes (not grandchildren classes) **leaf** – dumps data for leaf classes only (that is, child classes that don't have any children); this is the default setting for class data **all** – dumps data for the complete branch: the specified class, direct child classes, and all other classes that are descendants of the child classes. **Note:** The **immediate** and **leaf** options are not applicable to link data. If classes or partitions were deleted during this measurement interval, they will appear in the output list. Examples that compare **immediate**, **leaf**, and **all** options
<link-name>| *<partition-name>*| *<class-name>* *<host-ip-addr>*	List one or more names of the elements you want to dump; that is, specific class names, partition names, links (inbound or outbound), or host IP addresses. A specific name is not required if you want to dump *all* link, partition, class, or host data.
[endtime [*<date>*]*<time>*]	Specify a fixed end date and time for the data dump. The *<date>* should be in the formatYYYYMMDD and the *<time>* should be in the format HHMM. The *<date>* is optional, but the *<time>* is required. If **endtime** is omitted, the current date and time (now) are used.
<duration>	Time duration over which to total the values. Duration must be specified as one of: **s** for seconds, **m** for minutes, **h** for hours, **d** for days, **w** for weeks.
to *<file>*	Use the **to** literal to specify that the records should be dumped to a text file named *<file>*. The file will be saved in the current directory. If **to** *<file>* is omitted, the records are dumped to stdout. **Note:** This text file can be downloaded to a PC. You can then import the data into another program (such as Microsoft Excel) for further analysis.
[sort]	Use the **sort** literal to sort each dumped interval by the first variable, in descending order.
<var>...	Specify one or more variables to dump (for example, bytes, pkts, or avg-bps). To list the available variables, specify **?** for this parameter.

Examples:

```
measure dump class all inbound/http by element 1m peak-bps

 measure dump partition inbound/sales-dept by element 1d peak-bps partition-over-limit-
msecs

 measure dump class outbound/ftp by element 2d avg-bps class-hits policy-hits
```

Command Change History

Release	Modification
11.6.1	Partner option added

measure dump by event

List measurement data for adaptive response agents. For example, you can retrieve the score values and colors for each adaptive response agent in a certain time period.

measure dump [partner *<partner_ip>*] *<cdf-type>* all `[<agent-name>] [endtime [<date>]<time>]` **by event *<count>*|*<duration>*** `[to <file>] [sort]` ***<var>***...

where:

partner *<partner_ip>*	Dumps data for the specified standby partner. If a partner is not specified, data for the current PacketShaper is dumped.
<cdf-type>	The *<cdf-type>* can be one of the following: **agent1** retrieves scoring values and status color data for all agents **agent2** retrieves measurement data for High Bandwidth New App agent **agent3** retrieves measurement data for High Bandwidth Host agent **agent4** retrieves measurement data for Traffic Performance agents **agent5** retrieves measurement data for Partition Utilization agents
[*<agent-name>*]	Name of the agent, enclosed in quotes. For example: "Outbound Default Traffic" For a list of available agent names, use the **agent show** command.
[endtime [*<date>*]*<time>*]	Specify a fixed end date and time for the data dump. The *<date>* should be in the format YYYYMMDD and the *<time>* should be in the format HHMM. The *<date>* is optional, but the *<time>* is required. If **endtime** is omitted, the current date and time (now) are used.
to *<file>*	Use the **to** literal to specify that the records should be dumped to a text file named *<file>*. The file will be saved in the current directory. If **to***<file>* is omitted, the records are dumped to stdout. **Note:** This text file can be downloaded to a PC. You can then import the data into another program (such as Microsoft Excel) for further analysis.
[sort]	Use the **sort** literal to sort each dumped interval by the first variable, in descending order.
<count>	Retrieve data for a selected number of evaluation intervals
<duration>	Retrieve all data recorded over the selected duration. Use one of the following formats to indicate a time duration: *N*s for seconds, *N*m for minutes, *N*h for hours, *N*d for days, *N*w for weeks.

<var> The following measurement variables are available for each agent type:

agent1 variables: score-value, score-color, score-result

Note: Score colors are reported by the following values: 0=green, 1=red, 2=yellow, 3=blue. The score result has a value of 0 if the agent successfully measured its target. Otherwise, the score result will return an error code.

agent2 variables: element, sample-interval-msecs, namelist, avg-bps

agent3 variables: element, sample-interval-msecs, host-ip, direction, avg-bps. Direction is displayed as: 1=Inbound, 0=Outbound.

agent4 variables: element, sample-interval-msecs, class-id, network-efficiency

agent5: element, sample-interval-msecs, class-id, avg-bps

Examples:

The following command dumps data by *<count>*, and returns the data for the past two evaluation intervals for all agents (the ***agent1*** *<cdf-type>*).

me dump agent1 all by event 2 score-value score-colorscore-result

```
"time","agent1","score-value","score-color","score-result"
"13-Jan-2005 03:24:00","Class ME Variables agent",0,2,0
"13-Jan-2005 03:24:00","High Bandwidth Host",2,0,0
"13-Jan-2005 03:24:00","Inbound Default Traffic",1,0,0
"13-Jan-2005 03:24:00","Outbound Default Traffic",0,0,0
"13-Jan-2005 03:24:00","Partition Utilization agent",0,0,4557
"13-Jan-2005 03:24:00","Spoofing - Client",0,0,0
"13-Jan-2005 03:24:00","Spoofing - Server",0,0,0
"13-Jan-2005 03:24:00","Syn Attack - Failed Flows",0,0,0
"13-Jan-2005 03:24:00","Traffic Performance agent",2,1,0
"13-Jan-2005 03:23:00","Class ME Variables agent",0,2,0
"13-Jan-2005 03:23:00","High Bandwidth Host",2,0,0
"13-Jan-2005 03:23:00","Inbound Default Traffic",1,0,0
"13-Jan-2005 03:23:00","Outbound Default Traffic",0,0,0
"13-Jan-2005 03:23:00","Partition Utilization agent",0,0,4557
"13-Jan-2005 03:23:00","Spoofing - Client",0,0,0
"13-Jan-2005 03:23:00","Spoofing - Server",0,0,0
"13-Jan-2005 03:23:00","Syn Attack - Failed Flows",0,0,0
"13-Jan-2005 03:23:00","Traffic Performance agent",0,0,4557
```

To display data for a specific agent you can include the *<agent-name>*, enclosed in quotes:

measure dump agent1 all "Inbound Traffic Performance" by event 5 score-value

```
"agent1:Inbound Traffic Performance"
 "time","score-value"
 "12-Oct-2005 11:06:00",87
 "12-Oct-2005 11:05:00",100
 "12-Oct-2005 11:04:00",95
 "12-Oct-2005 11:03:01",92
 "12-Oct-2005 11:02:00",97
```

The command below returns data for a *<duration>*, displaying the last recorded minute of information for the **agent3** *<cdf-type>* (the High Bandwidth Host agent).

me dump agent3 all by event 1m host-ip direction avg-bps

```
"21-Jul-2005 13:51:00","High Bandwidth Host",174351134,0,56
"21-Jul-2005 13:51:00","High Bandwidth Host",2887062186,0,1824
"21-Jul-2005 13:51:00","High Bandwidth Host",2887061257,0,32
"21-Jul-2005 13:51:00","High Bandwidth Host",2887059946,0,8
"21-Jul-2005 13:51:00","High Bandwidth Host",2887057727,0,40
"21-Jul-2005 13:51:00","High Bandwidth Host",2887062278,0,16
"21-Jul-2005 13:51:00","High Bandwidth Host",2887061761,0,24
"21-Jul-2005 13:51:00","High Bandwidth Host",2887063396,0,16
"21-Jul-2005 13:51:00","High Bandwidth Host",2887063331,0,16
"21-Jul-2005 13:51:00","High Bandwidth Host",2887057710,1,224
"21-Jul-2005 13:51:00","High Bandwidth Host",2887062068,0,16
"21-Jul-2005 13:51:00","High Bandwidth Host",2887057428,0,344
"21-Jul-2005 13:51:00","High Bandwidth Host",2887059941,0,8
"21-Jul-2005 13:51:00","High Bandwidth Host",2887061002,0,8
"21-Jul-2005 13:51:00","High Bandwidth Host",2887059960,0,8
"21-Jul-2005 13:51:00","High Bandwidth Host",2887059943,0,8
"21-Jul-2005 13:51:00","High Bandwidth Host",2887060739,0,8
"21-Jul-2005 13:51:00","High Bandwidth Host",2887064948,0,24
"21-Jul-2005 13:51:00","High Bandwidth Host",2887064955,0,8
"21-Jul-2005 13:51:00","High Bandwidth Host",2887063300,0,16
"21-Jul-2005 13:51:00","High Bandwidth Host",2887062919,0,8
"21-Jul-2005 13:51:00","High Bandwidth Host",167837972,0,72
"21-Jul-2005 13:51:00","High Bandwidth Host",2887059250,1,120
```

Command Change History

Release	Modification
11.6.1	Partner option added
11.5.1	Command introduced

measure dump by time

Dump different time intervals of measurement data, with the most recent time listed first. As with the **by element** and **by var** options, **by time** lists measurement data totals for a specific element, such as a particular link (Inbound or Outbound), partition, class, or host IP address. You can list data for one or more measurement variables for a designated duration (such as 1 day or 2 weeks). Specific to the **by time** option, the amount of time in each interval must be specified; for example, you can break down the data into 10-minute intervals.

```
measure dump [partner <partner_ip>] link|partition|class|host [group day|month|accounting]
[immediate|leaf|all] [<link-name>|<partition-name>|<class-name>|<host-ip-addr>...]
[endtime [<date>]<time>] by time [element]all|<count>|<duration> <interval> [to <file>]
[sort] <var>...
```

partner *<partner_ip>*	Dumps data for the specified standby partner. If a partner is not specified, data for the current PacketShaper is dumped.
link|partition|class|host	Specify the type of data to be dumped.
[group day|month]	Dumps data from a specific measurement group day, month. If you don't specify a group, PacketShaper will choose the group that appears to most appropriate for the interval and/or duration you have specified. For example, if you specify an interval of 15 minutes, PacketShaper assumes you want to dump the data from the day group. Or if you specify an interval of 2 hours, PacketShaper assumes you want to dump the data from the month group. You can use the **group** option to override the default behavior.
[immediate|leaf|all]	**immediate**– dumps data for the specified class and direct child classes (not grandchildren classes) **leaf** – dumps data for leaf classes only (that is, child classes that don't have any children); this is the default setting for class data **all** – dumps data for the complete branch: the specified class, direct child classes, and all other classes that are descendants of the child classes. **Note:** The **immediate** and **leaf** options are not applicable to link data. If classes or partitions were deleted during this measurement interval, they will appear in the output list. Examples that compare **immediate**, **leaf**, and **all** options
<link-name>| *<partition-name>*| *<class-name>*| *<host-ip-addr>*	List one or more names of the elements you want to dump; that is, specific class names, partition names, links (inbound or outbound), or host IP addresses. A specific name is not required if you want to dump *all* link, partition, class, or host data.
[endtime [*<date>*]*<time>*]	Specify a fixed end date and time for the data dump. The *<date>* should be in the formatYYYYMMDD and the *<time>* should be in the format HHMM. The *<date>* is optional, but the *<time>* is required. If endtime is omitted, the current date and time (now) are used.

by time [element]	Use **time** to get a time-series data dump. Specify **element** to list the output as separate elements (links, partitions, classes, or hosts) within the time-series order. For example, when you specify **element** for a class, data for each class is output on a separate row.
	
	Example:
	
	`measure dump class all inbound by time element all 1d peak-bps`
	
	Result:
	
	"03-May-1998 15:27:02","/Inbound/Global/AppleTalk",0
	"03-May-1998 15:27:02","/Inbound/Global/NetBIOS",12858
	
	When you omit **element**, the output is displayed as a table, with column headers and a column for each element's data.
	
	Example:
	`measure dump class all inbound by time all 1d peak-bps`
	
	Result:
	
	"class"
	"time","/Inbound/Global/AppleTalk-peak-bps","/Inbound/Global/NetBIOS-peak-bps"
	"03-May-1998 15:27:02",0,25,12858
	"02-May-1998 15:54:35",0,32,98,11966
	"01-May-1998 15:15:33",0,4,2,15882
	"30-Apr-1998 15:30:30",0,0,0,23965

all|*<count>*|*<duration>*	Dump **all** available records for the specified interval; or dump the most recent *<count>* intervals; or dump all of the intervals within *<duration>* of the current time. Duration must be specified as one of: **s** for seconds, **m** for minutes, **h** for hours, **d** for days, **w** for weeks.
<interval>	Specify the interval time as one of: **s** for seconds, **m** for minutes, **h** for hours, **d** for days, **w** for weeks.
to *<file>*	Use the **to** literal to specify that the records should be dumped to the file named *<file>*. The file will be saved in the current directory. If **to** *<file>* is omitted, the records are dumped to stdout.
[sort]	Use the **sort** literal to sort each dumped interval by the first variable, in descending order. If sort is specified, a "by time element" dump format is used, even if a "by time" format was specified.
<var>...	Specify one or more variables to dump (for example, bytes, pkts, or avg-bps). To list the available variables, specify **?** for this parameter.

Examples

In the following example, the avg-bps and peak-bps variables are displayed for the Inbound partition. The duration is one week (1w) and the interval is one day (1d). In the output, each interval is a different record.

```
measure dump partition all inbound by time 1w 1d avg-bps peak-bps
```

"partition:/Inbound"
"time","avg-bps","peak-bps"
"09-Jun-2003 13:00:01",1707,2555221
"08-Jun-2003 13:00:01",2484,1796609

```
"07-Jun-2003 13:00:01",0,0
"06-Jun-2003 13:00:01",0,0
"05-Jun-2003 13:00:01",3436,8655317
"04-Jun-2003 13:00:01",983,2263233
"03-Jun-2003 13:00:01",4175,4219764
```

Here are some additional examples:

```
measure dump class inbound/global by time 5 1h avg-bps

 measure dump class all by time 5 1h avg-bps

 measure dump class all inbound/global by time element 5 1h avg-bps

 measure dump partition all inbound by time 1d 1h avg-bps peak-bps

measure dump class all endtime 200105261700 by time 5 1h bytes
```

To get data for just the /inbound and /outbound classes:

```
measure dump class immediate / by time 10m 1m bytes
```

Command Change History

Release	Modification
11.6.1	Partner option added

measure dump by var

List measurement data totals for a specific element, such as a particular link (Inbound or Outbound), partition, class, or host IP address. You can list data for one or more measurement variables for a designated duration (such as 1 day or 2 weeks). The measurement data can be displayed onscreen or stored in a delimited text file. The **by var** option dumps the same data as the **by element** option, except the output is structured differently. (examples that compare **by var** and **by element** output)

```
measure dump [partner <partner_ip>] link|partition|class|host [group day|month|accounting]
[immediate|leaf|all] [<link-name>|<partition-name>|<class-name>|<host-ip-addr>...] [endtime
[<date>]<time>] by var <duration> [to <file>] [sort] <var>...
```

partner *<partner_ip>*	Dumps data for the specified standby partner. If a partner is not specified, data for the current PacketShaper is dumped.
link|partition|class|host	Specify the type of data to be dumped. **Note:** In order to dump host measurement data, you must enable host accounting.

[group day|month|accounting]	Dumps data from a specific measurement group day, month, or accounting (if host accounting is enabled). If you don't specify a group, PacketShaper will choose the group that appears to most appropriate for the interval and/or duration you have specified. For example, if you specify an interval of 15 minutes, PacketShaper assumes you want to dump the data from the day group. Or if you specify an interval of 2 hours, PacketShaper assumes you want to dump the data from the month group. You can use the **group** option to override the default behavior.
[immediate|leaf|all]	**immediate**– dumps data for the specified class and direct child classes (not grandchildren classes) **leaf** – dumps data for leaf classes only (that is, child classes that don't have any children); this is the default setting for class data **all** – dumps data for the complete branch: the specified class, direct child classes, and all other classes that are descendants of the child classes. **Note:** The **immediate** and **leaf** options are not applicable to link data. If classes or partitions were deleted during this measurement interval, they will appear in the output list. Examples that compare **immediate**, **leaf**, and **all** options
<link-name>| <partition-name>| <class-name>| <host-ip-addr>	List one or more names of the elements you want to dump; that is, specific class names, partition names, links (inbound or outbound), or host IP addresses. A specific name is not required if you want to dump *all* link, partition, class, or host data. **Note:** As mentioned above, host accounting must be enabled in order to get measurement data on specific hosts.
[endtime [*<date>*]*<time>*]	Specify a fixed end date and time for the data dump. The *<date>* should be in the formatYYYYMMDD and the *<time>* should be in the format HHMM. The *<date>* is optional, but the *<time>* is required. If endtime is omitted, the current date and time (now) are used.
<duration>	Time duration over which to total the values. Use one of the following: **s** for seconds, **m** for minutes, **h** for hours, **d** for days, **w** for weeks.
to *<file>*	Use the **to** literal to specify that the records should be dumped to the file named *<file>*. The file will be saved in the current directory. If **to***<file>* is omitted, the records are dumped to stdout.
[sort]	Use the **sort** literal to sort each dumped interval by the first variable, in descending order. If sort is specified, a "by time element" dump format is used, even if a "by time" format was specified.
<var>...	Specify one or more variables to dump (for example, bytes, pkts, or avg-bps). To list the available variables, specify **?** for this parameter. When listing host accounting data, you can specify bytes, kbytes, or host accounting categories for the *<var>*.

Examples:

```
measure dump partition all /inbound/http by var 8h avg-bps peak-bps

 measure dump class /inbound/http by var 1d to daily.out avg-bps

 measure dump class all /inbound/http by var 1w sort peak-bps pkts
```

"time:17-Jun-1998 18:21:48"
"class-var","/Inbound/HTTP/Default","/Inbound/HTTP/gifs"
"peak-bps",229729,86944
"pkts",39,26

An example that lists data for a standby partner:

```
measure dump partner 10.10.10.10 class all /inbound/http by var 2h sort peak-bps
```

An example that lists host accounting data for a specific host:

```
measure dump host 152.163.209.65 by var 30m sort bytes web mail
```

"time:08-May-2002 15:20:00"
"host-var","152.163.209.65"
"bytes",2730
"web",2316
"mail",0

Command Change History

Release	Modification
11.6.1	Partner option added

Measure Dump Examples

Here is the portion of the traffic tree that is being dumped:

```
Inbound
 HTTP
 CriticalWeb
 test
 Default
 OtherWeb
```

To see the difference between the three options (**immediate**, **leaf**, and **all**), look at the output in the following examples.

Immediate option

me dump class immediate inbound/http by var 1d bytes avg-bps

"time:09-Jul-2002 15:46:00"
"class-var","/Inbound/HTTP","/Inbound/HTTP/CriticalWeb","/Inbound/HTTP/OtherWeb"
"bytes",419859,0,419859
"avg-bps",133,0,133

Leaf option

me dump class leaf inbound/http by var 1d bytes avg-bps

"time:09-Jul-2002 15:47:00"
"class-var","/Inbound/HTTP/CriticalWeb/test"," /Inbound/HTTP/CriticalWeb/Default"," /Inbound/HTTP/OtherWeb"
"bytes",0,0,419859
"avg-bps",0,0,133

All option

me dump class all inbound/http by var 1d bytes avg-bps

"time:09-Jul-2002 15:48:00"
"class-var","/Inbound/HTTP","/Inbound/HTTP/CriticalWeb"," /Inbound/HTTP/CriticalWeb/test","
/Inbound/HTTP/CriticalWeb/Default","/Inbound/HTTP/OtherWeb"
"bytes",419859,0,0,0,419859
"avg-bps",133,0,0,0,133

Measure Dump "by var" and "by element" Comparison

The **by var** option dumps the same measurement data as the **by element** option, except the output is structured differently. In the **by var** output, the data associated with each *measurement variable* is listed in a separate record (row). In the **by element** output, the data associated with each *element* (link, partition, class, or host) is listed in a separate record. To see the difference, compare the output in the following examples.

"by var" option

The following example requests the data, for an hour duration, recorded for three measurement variables (peak-bps, avg-bps, and bytes) for the Inbound/DiscoveredPorts classes. The **by var** option is used so that each variable's data is in a separate row.

measure dump class all inbound/discoveredports endtime 200306090900 by var 1h peak-bps avg-bps bytes

"time:09-Jun-2003 07:00:02"
"class-var","/Inbound/DiscoveredPorts","/Inbound/DiscoveredPorts/UDP_Port_1103","/Inbound/DiscoveredPorts/UDP_
Port_1108","/Inbound/
DiscoveredPorts/UDP_Port_2222","/Inbound/DiscoveredPorts/UDP_Port_2930","/Inbound/DiscoveredPorts/UDP_Port_
4734","/Inbound/
DiscoveredPorts/UDP_Port_5165","/Inbound/DiscoveredPorts/UDP_Port_5353","/Inbound/DiscoveredPorts/UDP_Port_
7339","/Inbound/
DiscoveredPorts/UDP_Port_9759"
"peak-bps",752,0,0,0,0,0,0,752,0,0
"avg-bps",97,0,0,0,0,0,0,97,0,0
"bytes",43492,0,0,0,0,0,0,43492,0,0

If this data were to be imported into Microsoft Excel, each class name (such as /Inbound/DiscoveredPorts) would be a column heading and each variable name (such as peak-bps) would be a row heading.

"by element" option

This example also requests the data, for an hour duration, recorded for three measurement variables (peak-bps, avg-bps, and bytes) for the Inbound/DiscoveredPorts classes. The **by element** option, though, places the data for each element (class, in this case) in a separate record.

measure dump class all inbound/discoveredports endtime 200306090900 by element 1h peak-bps avg-bps bytes

"time:09-Jun-2003 07:00:02"
"class","peak-bps","avg-bps","bytes"
"/Inbound/DiscoveredPorts",752,97,43492
"/Inbound/DiscoveredPorts/UDP_Port_1103",0,0,0
"/Inbound/DiscoveredPorts/UDP_Port_1108",0,0,0
"/Inbound/DiscoveredPorts/UDP_Port_2222",0,0,0
"/Inbound/DiscoveredPorts/UDP_Port_2930",0,0,0
"/Inbound/DiscoveredPorts/UDP_Port_4734",0,0,0
"/Inbound/DiscoveredPorts/UDP_Port_5165",0,0,0
"/Inbound/DiscoveredPorts/UDP_Port_5353",752,97,43492
"/Inbound/DiscoveredPorts/UDP_Port_7339",0,0,0
"/Inbound/DiscoveredPorts/UDP_Port_9759",0,0,0

If this data were to be imported into Microsoft Excel, each variable name (such as peak-bps) would be a column heading and each class name (such as /Inbound/DiscoveredPorts) would be a row heading.

measure reset

Reset the measurement configuration to its factory-default state. This command clears the measurement data stored on the unit s 9.258/me/measure folder. Resetting may be necessary after upgrading the software. (See the last paragraph below for further explanation.)

```
measure reset [link|partition|class|host]
```

With the optional **[link|partition|class|host]** parameters, you can selectively clear different types of measurement data.

If a parameter is not specified, all types of measurement data are cleared.

After you issue the command, you are prompted to confirm your reset request. Accumulated measurement data is cleared and the unit resets. Measurement and reporting data will not be available for several minutes.

After **measure reset** executes, the measurement engine begins running as a background process. During the first ten minutes following the reset, do not attempt to load a new software image, as the file transfer will conflict with measurement engine processing.

You will need to reset measurement data in order to use new measurement variables introduced in new PacketShaper software versions. To help you determine whether you need to reset measurement data, the output from the **measure show** command

indicates if the measurement reset has been done. For example, if the message A Measurement Reset of this type has not been done appears in the **measure show class** output, PacketShaper has detected that you upgraded to an image that has new measurement variables of the type you listed. To enable the new variables of that type, you need to issue the **measure reset class** command.

measure show

Use the **measure show** command to check the measurement engine status, to see whether the measurement engine needs to be reset, or to display the details for a specific measurement type or standby partner.

```
measure show [partner <partner_ip>
```

The resulting display indicates the status of the current PacketShaper's (or standby partner's) measurement engine: running, starting, or stopped. The output displays one line of status for each element in the measurement volume. Each line contains the element type, the sample interval (m=minutes and h=hours), the number of recorded samples, and the number of samples that can be recorded.

Measurement engine is running
A complete Measurement Reset has not been done.

```
link day        Interval:1m Samples:   20160/   20160
class day       Interval:1m Samples: 4042858/14414400
class month     Interval:1h Samples:   67337/ 7447440
partition month Interval:1h Samples:    1802/ 7447440
partition day   Interval:1m Samples:  108503/14414400
link month      Interval:1h Samples:    1802/    8928
```

If the message "A complete Measurement Reset has not been done" appears in the **measure show** output (as shown above), PacketShaper has detected that you upgraded to an image that has new measurement variables, or that you loaded a license that affects the number of classes. To reset the measurement database, you need to issue the measure reset command.

To display the measurement engine status for a specific standby partner, you can enter its IP address.

```
measure show partner 10.9.64.155
Measurement engine is running

Standby partner 0 [0314320399]
partition day   Interval:1m Samples:   97312/14414400
class month     Interval:1h Samples:  171295/ 7447440
link day        Interval:1m Samples:   20160/   20160
link month      Interval:1h Samples:    1802/    8928
class day       Interval:1m Samples:10286539/14414400
partition month Interval:1h Samples:    1620/ 7447440
```

Displaying Measurement Element Status

To display the status of a particular element type, use:

```
measure show [partner <partner_ip>] link|partition|class|host
```

Example:

```
measure show class
```

Measurement engine is running
class base interval: 60 (secs)

A Measurement Reset of this type has not been done.

class day Interval: 1m Samples: 438/ 368640
class month Interval: 1h Samples: 0/ 190464

If the message "A Measurement Reset of this type has not been done" appears in the **measure show** output (as shown above), PacketShaper has detected that you upgraded to an image that has new measurement variables, or that you loaded a license that affects the number of classes. To reset the measurement database, you need to issue the **measure reset** command.

Listing Measurement Variables

Display the status and available variables for a particular measurement group.

```
measure show [partner <partner_ip>] [link|partition|class|host] [day|month]
measure show host accounting
```

Examples:

```
measure show class day
```

```
Group:          class day
Interval time:  60 (seconds)
Interval count: 7194 (minimum)
Path:           /opt/bluecoat/ps/measure/cl_day.dat
Samples:        max 72014400 current 2936938
Overruns:       1386953639
Last sample:    at 13-Dec-2013 11:50:00
Next sample:    13-Dec-2013 11:51:00 to near 13-Dec-2013 11:52:00
Variable(s):    element sample-interval-msecs sample-interval-overruns bytes
                pkts tcp-data-pkts tcp-retx-pkts tcp-early-retx-toss-pkts
                guar-rate-fails guar-rate-allocs peak-tcp-conns tcp-conn-inits
                tcp-conn-exits tcp-conn-server-refuses tcp-conn-server-ignores
                tcp-conn-aborts tcp-conn-self-denies rtp-expected-pkts
                rtp-lost-pkts rtp-jitter-count rtp-jitter-sum
                peak-guar-rate-flows tcp-retx-bytes
                ipv6-bytes ipv4-bytes noip-bytes class-hits policy-hits
                conn-speed-hist peak-bps license-overflows licenses-total
                licenses-peak cfg-agg-threshold cfg-good-transaction
                service-level-errors total-trans trans-bytes
                client-flood-block peak-ipdg-conns web-response-2XX
                web-response-3XX web-response-4XX web-response-5XX
                rtp-latency-sum rtp-latency-count icmp-rtt-msecs-sum
```

```
icmp-rtt-count no-bw-drop-bytes total-bytes-tx is-rtp
icmp-attempt-count drop-bytes% icmp-rtt-msecs-avg
icmp-success-rate% rtp-lost-pkt% rtp-jitter-avg rtp-latency-avg
rfactor mos-x10 sample-interval-secs kbytes avg-bps avg-pps
tcp-retx-pkts% tcp-efficiency% tcp-early-retx-toss-pkts%
tcp-conn-server-refuses% tcp-conn-server-ignores%
tcp-conn-aborts% tcp-conn-self-denies% pkt-exchange-time-avg
app-availability% server-flood-block service-level%
trans-bytes-avg avg-round-trip-time
service-level-threshold%
```

Overruns (listed in the output above) occur when the measurement daemon misses a full (1-minute) interval.

Command Change History

Release	Modification
11.6.1	Partner option added

measure start

Restart the measurement engine after it has been stopped with the measure stop command.

```
measure start
```

After you issue the **measure start** command, PacketShaper will begin recording measurement data (although not for the interval in which it was started).

measure stop

Stop or pause the measurement engine.

```
measure stop
```

After you issue this command, PacketShaper will stop recording measurement data until you restart the measurement engine with the measure start command. Data will not be recorded for the interval in which it was stopped.

> **Caution:** Use the **measure stop** command with caution. When the measurement engine is stopped, no data will be recorded. Some features, such as Top Ten, will not function properly without current data.

mib

In addition to the SNMP Management Information Base (MIB) files, PacketShaper supports a variety of other internal MIBs. These internal MIBs contain data such as MAC cache and DNS information. These diagnostic commands are intended to be used only under the guidance of Customer Support and are not covered in this guide.

more

Display the named file, showing a single page and pausing before displaying the next page. More than one filename can be specified.

```
more [-<number>] <filename>
```

Providing an optional number will display the specified number of lines on one page.

net nic

View network statistics such as packets transmitted and discarded. These statistics are accumulated since the last PacketShaper reset.

```
net [nic {<device>}]|ip|pna
```

nic	Show Ethernet statistics. The parameters inside and outside represent ports 0 and 1, respectively. To display both inside and outside statistics, omit this parameter.
ip	Show IP statistics
pna	Show network statistics

where *<device>* is the interface name:

Dual-port NICs:

- slot1_in
- slot1_out
- slot2_in
- slot2_out
- slot3_in
- slot3_out
- slot4_in
- slot4_out

Quad-port NICs:

- slot3_in1
- slot3_out1
- slot3_in2
- slot3_out2
- slot4_in1
- slot4_out1
- slot4_in2
- slot4_out2

Command Change History

Release	Modification
11.1.1	Updated to support PS-Series device names

packetcapture add

Specify the type of traffic for which you want to capture packets. The packet capture feature can capture packets for future analysis, allowing you to analyze detailed information about the packets, such as the source and destination IP addresses and protocols used. You can capture packets for traffic classes, IP addresses and ranges, subnets, host lists, port numbers and ranges, and NICs.

```
packetcapture add class:<tclass> | list:<hostlist> | host:<ipaddr> | net:<ipaddr>/<cidr> |
range:<low>-<high> | port:<low>-<high> | nic:<interface pair> [<ipaddr>|<ipaddr>/<cidr>]
```

class: *<tclass>*

Name of the traffic class. Specify the complete path if the class exists in both Inbound and Outbound. The class can be based on any protocol, including IPv6.

You can capture packets from multiple classes, but you must add them one class at a time.

Packet capture can only capture packets for leaf classes (classes without any children). For example, if your Inbound/HTTP class has child classes (such as /Inbound/HTTP/Critical), you cannot capture packets on /Inbound/HTTP.

Example: *class:inbound/default*

list: *<hostlist>*

Name of the host list created with the hl new command

Example: *list:finance*

host: *<ipaddr>*

IP address or domain name of an IPv4 host; IPv6 hosts are not supported.

Examples: *host:172.21.18.160*, *host:west.us.com*

net: *<ipaddr>*/*<cidr>*

The address of the IPv4 subnet; the CIDR number specifies the number of constant bits in the address range. IPv6 subnets are not supported.

Example: *net:10.0.0.0/8*

range: *<low>-<high>*	Range of IPv4 addresses, separated by a dash Example: *range:192.21.18.160-192.21.18.170*
port: *<low>-<high>*	Range of port numbers, separated by a dash Examples: *port:80*, *port:1000-3000*
nic: *<interface pair>* [*<ipaddr>*/*<ipaddr>*/*<cidr>*]	Network Interface Card (NIC) where *<interface pair>* is the device name, such as *slot1* or *slot3_pair1* Example: *nic:slot1* Optionally, a host or a subnet can be specified for the NIC pair. Example: *nic:slot4_pair2 10.0.0.0/8*

Additional Information

- When the NIC filter is used, all other filters will be removed.
- Two output files will be created: one for packets captured as they are received by the NIC's Inside port (xxxxxx**in**.dmp) and one for packets captured as they leave the NIC on its Output port (xxxxxx**ot**.dmp). The *xxxxxx* part of the filename specifies the date and time of the capture file.
- The packetcapture limit packets command does not apply to NIC filters since this filter captures raw packets from the NIC card without any processing by the PacketShaper.

Packet capturing doesn't begin until after the feature is enabled (see packetcapture on). When you no longer want to capture packets for a filter you have added, use the packetcapture remove command.

Command Change History

Release	Modification
11.1.1	Updated to support PS-Series device names
8.4.1	Change to NIC filter: traffic on the management port can be captured
8.2.2, 8.3.1	Added option to filter by host or subnet at the NIC level
8.2.0	NIC filter added
8.1.1	Additional filters for IP addresses and ranges, subnets, host lists, ports

packetcapture limit packets

Set the number of packets (1 255) that will be captured per flow. If you don t specify a limit, an unlimited number of packets will be captured. By limiting the number of packets per flow, you will be able to capture more flows. Note that the number of packets per flow is for both directions combined, so if both directions are being logged, the number of packets is the sum of the two directions.

```
packetcapture limit packets <n>
```

where *<n>* is the number of packets per flow. Typical limit values are 10 or 20.

If you want to remove the limit, use:

```
packetcapture limit packets none
```

> **Note:** The **limit packets** option does not apply to NIC filters.

packetcapture off

Disable packet capture and close the log file.

```
packetcapture off
```

After packet capture is turned off, you can download the file to a local workstation and use a protocol analyzer such as EtherPeek to look at the contents of the file. Use the packetcapture status command to determine the filename to download.

packetcapture on

Enable packet capture. After capture is turned on, PacketShaper will create a .dmp file in TCPDump format and start capturing packets into this file.

The **packetcapture on** command checks to verify that there is enough space on the PacketShaper before launching the packet logger utility. If there is insufficient space on the PacketShaper to write a full packet capture dump file, the PacketShaper CLI displays a warning that there is not enough memory to activate the feature. Remove any unnecessary files from the PacketShaper before reissuing the command.

> **Note:** Before you can enable packet capture, you must add at least one item to the capture list (see packetcapture add). You may also want to limit the number of packets that are captured per flow (see packetcapture limit packets).

```
packetcapture on
```

The file is stored in the 9.258/pktlog folder and is named according to the day and time packet capture was enabled. For example, if capture was enabled on the 12th of the month at 9:02:35am, the filename would be 12090235.dmp. When NIC filters are used, two output files are created: one for packets captured as they are received by the NIC's Inside port (xxxxxx**in**.dmp) and one for packets captured as they leave the NIC on its Output port (xxxxxx**ot**.dmp). The *xxxxxx* part of the filename specifies the date and time of the capture file.

Packets are captured until one of the following occurs:

- capturing is disabled with the packetcapture off command
- the log file reaches 99% of the maximum log file size (use the packetcapture status command to see the current and maximum log file sizes; these sizes vary by model)

packetcapture remove

Remove a filter that was added to the packet capture list (using the packetcapture add command).

```
packetcapture remove class:<tclass> | list:<hostlist> | host:<ipaddr> | net:<ipaddr>/<cidr>
| range:<low>-<high> | port:<low>-<high> | nic:<interface pair>
```

class: *<tclass>*	Name of the traffic class. Specify the complete path if the class exists in both Inbound and Outbound. Example: *class:inbound/default*
list: *<hostlist>*	Name of the host list Example: *list:finance*
host: *<ipaddr>*	IP address or domain name of a host Examples: *host:172.21.18.160*, *host:west.us.com*
net: *<ipaddr>*/*<cidr>*	The address of the subnet; the CIDR number specifies the number of constant bits in the address range Example: *net:10.0.0.0/8*
range: *<low>*-*<high>*	Range of IP addresses, separated by a dash Example: *range:192.21.18.160-192.21.18.170*
port: *<low>*-*<high>*	Range of port numbers, separated by a dash Examples: *port:80*, *port:1000-3000*

nic:*<interface pair>*	Network Interface Card (NIC) where *<interface pair>* is one of the following: Dual-port NICs: ■ slot1_in ■ slot1_out ■ slot2_in ■ slot2_out ■ slot3_in ■ slot3_out ■ slot4_in ■ slot4_out Quad-port NICs: ■ slot3_in1 ■ slot3_out1 ■ slot3_in2 ■ slot3_out2 ■ slot4_in1 ■ slot4_out1 ■ slot4_in2 ■ slot4_out2 Management interface: ■ mgmt Example: *nic:slot1_in*

To see a list of the current packet capture filters, use the packetcapture status command.

Command Change History

Release	Modification
11.1.1	Updated to support PS-Series device names
8.2.0	NIC filter added
8.1.1	Additional filters for IP addresses and ranges, subnets, host lists, ports

packetcapture status

List the current packet capture settings. The status report indicates whether packet capture is enabled, the name and location of the log file, the file format, the maximum and current size of the log file, the number of packets in the current log file, and which items are being logged.

```
packetcapture status
```

Example output:

```
  Packet capture status:    OK
  Packet capture:           On - Logging
  Log file directory:       9.258/pktlog
  Log file name:            15153739.dmp
  Log file format:          tcpdump
  Maximum log size:         1006632960  bytes
  Current log size:         640 bytes (0%)
  Packets in current log:   8
  Captured option(s):
                            IP Host: 10.1.1.70 Subnet mask 255.255.255.255
                            class:/Inbound/Default
                            Port Range: 80 - 80
                            IP Range: 192.21.18.160-192.31.18.170
```

Additional Information

- The maximum log size on the PS-S500 is 1 GB for each direction. If you have multiple pcap files, it can be limited by the available disk space.
- When a NIC filter is configured, the **packetcapture status** output will show the status of two output files: one for the Inside port and one for the Outside.

partition apply

Create a static partition for a traffic class.

```
partition apply <tclass><minBps>|<minPct>|uncommitted [<maxBps>|<maxPct>|none|fixed]
```

<tclass>	This traffic class and all of its children are partitioned together (those that are not already separately partitioned), so this traffic class becomes the root of a partitioned subtree of traffic classes.

<minBps> | *<minPct>* | uncommitted	The minimum size of the new partition, specified in bits per second (*minBps*) or as a percentage of the parent partition's minimum size (*minPct*). If *<minPct>* is used, you must include the percent sign (for example, 10%). The minimum partition size is 1000 bps. The sum of all the partitions within either Inbound or Outbound can exceed the link size, allowing you to oversubscribe the link. Use the literal **uncommitted** to indicate that the guaranteed minimum allocation is whatever is not committed to other partitions. Normally uncommitted is used only by the default Inbound and Outbound partitions.
[*<maxBps>* | *<maxPct>* | none | fixed]	Limit the maximum bandwidth used by a burstable partition. The maximum can be specified in bits per second (*maxBps*) or as a percentage of the parent maximum (*maxPct*). If *<maxPct>* is used, you must include the percent sign (for example, 10%). The maximum must be greater than the minimum. The maximum partition size varies by model. PS-S200: licensed bandwidth PS-S400: 2 Gbps or licensed bandwidth (whichever is less) PS-S500: 2 Gbps. Exception: The Inbound and Outbound partitions can have a maximum size of 10 Gbps if you have a 10G license key or 5 Gbps if you have a 5G key. ■ Specify *<maxBps>* or *<maxPct>* to allow the partition to burst up to the specified bandwidth. ■ Specify **none** to allow the partition to burst to any available bandwidth. ■ Specify **fixed** to not burst, i.e., prevent a partition from exceeding the *<minBps>* or *<minPct>* size. If you do not specify anything for [*<maxBps>* | *<maxPct>* | none | fixed], the partition defaults to **none** (burstable to any available bandwidth).

Additional Information

- In order for partitions to take effect, traffic shaping must be enabled. See setup shaping.
- When creating partitions, make sure you don't allocate bandwidth in such a way that Inbound/Default and Outbound/Default get "starved," that is, there is no bandwidth available for these classes. If this happens, traffic classification and policies may not work as expected.

Examples:

Create an inbound burstable partition (no maximum limit specified) of 10000 bps:

```
partition apply inbound/outside/http 10k
```

Create an inbound burstable partition of 20000 bps with the ability to borrow additional bandwidth from other partitions, if it is available, up to 30000 bps:

```
partition apply inbound/outside/http 20k 30k
```

Create a burstable partition for SAP that is 30% of the link (Inbound partition) size, with a maximum size of 40%:

```
partition apply inbound/sap 30% 40%
```

In the above example, if the link size is 1.5 Mbps, the SAP partition would get a minimum of 450 Kbps and a maximum of 600 Kbps.

partition dynamic apply

Create a dynamic per-user partition for a traffic class. Can be specified by IP address or by subnet.

> **Note:** Before you can create dynamic per-user partitions, you must create a static partition for the class using the **partition apply** command.

To create a dynamic partition which handles traffic for an IP address:

`partition dynamic apply <tclass> per-address <side> <minBps>|<minPct>|uncommitted <maxBps>|<maxPct>|none|fixed`

<tclass>	Name of the traffic class having a static partition you would like to subdivide for each user
<side>	Side (**inside** or **outside**) of the PacketShaper on which the user is located
<minBps> | *<minPct>* | uncommitted	Minimum amount of bandwidth to be assigned to each user, specified in bits per second (*minBps*) or as a percentage of the parent partition s size (*minPct*). If *<minPct>* is used, you must include the percent sign (for example, 10%). The minimum subpartition size is 1000 bps. Use the literal **uncommitted** to indicate that the guaranteed minimum allocation is whatever is not committed to other partitions. Set this field to zero (0) to have PacketShaper allocate bandwidth equitably to each subpartition, so that the total of all subpartitions equals the static partition's size. **Note:** Minimum subpartition size is usually best handled by setting this field to zero and setting a maximum number of subpartitions (using the partition dynamic cap command). However, you must use a non-zero size if you want to implement per-session guaranteed rates within rate policies for this same traffic.
<maxBps> | *<maxPct>* | none | fixed	Maximum amount of bandwidth to be assigned to each subpartition, specified in bits per second (*maxBps*) or as a percentage of the parent partition s size (*maxPct*). If *<maxPct>* is used, you must include the percent sign (for example, 10%). Specify a *<maxBps>* value if you want to enforce a cap on each user or subnet even if more bandwidth is available. Managed bandwidth service providers are most frequently in this position, needing to cut off usage at agreed-upon, paid-for limits. If you don t want a maximum, specify **none**. Specify **fixed** to prevent a subpartition from exceeding the *<minBps>* or *<minPct>* size. Even if this field is left blank, the limit on the static, parent partition still restricts the total bandwidth for the aggregate of all subpartitions.

To create a dynamic partition which handles traffic for a subnet:

```
partition dynamic apply <tclass> per-subnet /<cidr> <side><minBps>|<minPct>
<maxBps>|<maxPct>|none
```

<tclass>	Name of the traffic class having a static partition you would like to subdivide for each user
/*<cidr>*	CIDR number specifying the number of constant bits in the address range
<side>	Side (**inside** or **outside**) of the PacketShaper on which the user is located
<minBps> | *<minPct>* | uncommitted	Minimum amount of bandwidth to be assigned to each user, specified in bits per second (*minBps*) or as a percentage of the parent partition s size (*minPct*). If *<minPct>* is used, you must include the percent sign (for example, 10%). The minimum sub partition size is 1024 bps. Use the literal **uncommitted** to indicate that the guaranteed minimum allocation is whatever is not committed to other partitions. Set this field to zero (0) to have PacketShaper allocate bandwidth equitably to each subpartition, so that the total of all subpartitions equals the static partition's size. **Note:** Minimum subpartition size is usually best handled by setting this field to zero and setting a maximum number of subpartitions (using the partition dynamic cap command). However, you must use a non-zero size if you want to implement per-session guaranteed rates within rate policies for this same traffic.
<maxBps> | *<maxPct>* | none | fixed	Maximum amount of bandwidth to be assigned to each subpartition, specified in bits per second (*maxBps*) or as a percentage of the parent partition s size (*maxPct*). If *<maxPct>* is used, you must include the percent sign (for example, 10%). Specify a maximum value if you want to enforce a cap on each user or subnet even if more bandwidth is available. Managed bandwidth service providers are most frequently in this position, needing to cut off usage at agreed-upon, paid-for limits. If you don t want a maximum, specify **none**. Specify **fixed** to prevent a subpartition from exceeding the *<minBps>* or *<minPct>* size. Even if this field is left blank, the limit on the static, parent partition still restricts the total bandwidth for the aggregate of all subpartitions.

After the dynamic partition is set up, whenever a new user begins generating flows in that class, a subpartition will be created for the user on the fly. The per-user partition remains in existence until it's re-used for new flows by the same user or needed by another user. A subpartition may be given to another user if there have not been any recent flows in the partition. To be more precise, a subpartition may be given to another user if 30 seconds have passed without any flows or if it's been five minutes since an established flow has sent any packets.

partition dynamic cap

Set the maximum number of active users allowed in the dynamic partition, and, optionally, create an overflow partition for users to tap into if the cap is exceeded.

```
partition dynamic cap <tclass><maxusers> [<overflowMinBps>|<overflowMinPct>|uncommitted
<overflowMaxBps>|<overflowMaxPct>|none|fixed]
```

<tclass>	Name of the traffic class having a dynamic partition for which you would like to set a cap
<maxusers>	Maximum number of per-user partitions that can be created in this traffic class

<overflowMinBps> | *<overflowMinPct>* | uncommitted	Minimum amount of bandwidth in the overflow partition, specified in bits per second (*overflowMinBps*) or as a percentage of the parent partition s size (*overflowMinPct*). If *<overflowMinPct>* is used, you must include the percent sign (for example, 10%). The minimum overflow subpartition size is 1024 bps. Use the literal **uncommitted** to indicate that the guaranteed minimum allocation is whatever is not committed to other partitions.
<overflowMaxBps> | *<overflowMaxPct>* | none | fixed	Maximum amount of bandwidth in the overflow partition, specified in bits per second (*overflowMaxBps*) or as a percentage of the parent partition s size (*overflowMaxPct*). If *<overflowMaxPct>* is used, you must include the percent sign (for example, 10%). When a value is specified, the overflow partition can use available excess bandwidth if needed. Specify **none** to allow the overflow partition to use any available bandwidth. Specify **fixed** to prevent the partition from exceeding the *<overflowMinBps>* or *<overflowMinPct>* specification. If you don t specify a value, the overflow partition has a fixed size; when it s not using its reserved bandwidth, that bandwidth is available to other traffic.

To remove the cap on a dynamic partition:

```
partition dynamic cap <tclass> none
```

partition dynamic cleanup

Remove idle dynamic partitions. Use this command if you don't want to wait for the automatic cleanup process that occurs every minute.

```
partition dynamic cleanup
```

The output indicates the number of dynamic partitions that were removed. For example:

```
Cleanup 21 dynamic partitions
```

> **Note:** By default, active dynamic partitions are retained for 5 minutes and idle dynamic partitions are retained for 30 seconds. System variables are available to change these default settings. See setup variable.

partition dynamic remove

Remove a dynamic partition. The partition reverts to being static.

```
partition dynamic remove <tclass>
```

partition dynamic summary

Show all the configured dynamic partitions and the number of users currently using each partition. Users that are not active can be replaced by new users.

```
partition dynamic summary
```

Example:

```
partition dynamic summary
```

Partition Name	--- Users --- Current	Cap	--- Current User Details --- Active	Idle	Gone	LongGone
Inbound http	7	none	3	0	3	1

After a dynamic partition is set up, whenever a new user begins generating flows in that class, a subpartition will be created for the user on the fly. The per-user partition remains in existence until it's re-used for new flows by the same user or needed by another user. A subpartition may be given to another user if there have not been any recent flows in the partition.

A subpartition is considered **Idle** if it has not been active for 300 seconds (5 minutes). Idle subpartitions still have flows which are sending packets. A subpartition is considered **Gone** if the flows associated with it have been gone 30 seconds or less, or **LongGone** if they have been gone more than 30 seconds. When the dynamic partition cap has been reached, new subpartitions are created from LongGone and Gone partitions.

In other words, a subpartition may be given to another user if 30 seconds have passed without any flows or if it's been five minutes since an established flow has sent any packets.

partition remove

Remove a static partition from a traffic class. The bandwidth allocated to this traffic class is returned to the parent partition.

```
partition remove <tclass>
```

partition show

Display current partition usage for static, dynamic, or both types of partitions.

```
partition show [<tclass>] [static|dynamic {<ip-addr>|<subnet>/<cidr>|<ip-addr-range>}|clear|config]
```

where:

<tclass>	Displays partition statistics for the specified traffic class
static	Lists only static partitions
dynamic	Lists only dynamic partitions
dynamic {*<ip-addr>*|*<subnet>*/*<cidr>*|*<ip-addr-range>*}	Lists only dynamic partitions for the specified IP address, subnet, CIDR, or range of IP addresses
clear	Resets the displayed partition statistics and then displays a partition list. Note that the statistics will not necessarily show as zero after this reset, because traffic activity could be recorded instantaneously. **Note:** The **clear** option is not applicable in conjunction with the ***<tclass>*** option.
config	Displays the minimum and maximum usage. The usage maximum is a partition's burst limit.

Examples:

To display partition statistics for all partitions, omit all parameters.

```
PacketShaper# partition show

Partition name          Size          Grntd       Prior       Curr  1-Min
                        Min /  Max          Excess      Usage       Avg   Peak
-------------------------------------------------------------------------------
/Inbound                 1.5M   1.5M     0     0  1024   1024  2048  3687  537k
  /Inbound/MPEG-Audio    500k   1.5M*    0     0     0      0     0     0    0
  /Inbound/WinMedia        0    1.5M*    0     0     0      0     0     0    0
-------------------------------------------------------------------------------
/Outbound                1.5M   1.5M     0     0     0     0  1143   405  6986
/Outbound                1.5M   1.5M     0     0     0     0  1143   405  6986
 /Outbound/157             0    1.5M     0     0     0     0     0     0     0
 /Outbound/157             0    1.5M     0     0     0     0     0     0     0
  /Outbound/157/74.125.77.99
                           0    100k     0     0     0     0     0     0     0
  /Outbound/157/74.125.77.102
                           0    100k     0     0     0     0     0     0     0
  /Outbound/157/10.2.2.100
                           0    100k     0     0     0     0     0     0     0
  /Outbound/157/91.189.90.41
                           0    100k     0     0     0     0     0     0     0
  /Outbound/157/74.125.77.104
                           0    100k     0     0     0     0     0     0     0
  /Outbound/157/74.125.79.102
                           0    100k     0     0     0     0     0     0     0
  /Outbound/157/91.189.90.40
                           0    100k     0     0     0     0     0     0     0
```

This output lists both minimum and maximum partition size settings. It also lists the rate of priority traffic. In addition, it displays an asterisk (*) next to any minimum or maximum value that isn't "pure" that is, if the programmed value was adjusted due to (1)

oversubscription or (2) the use of the strings **fixed** or **none**. The adjusted values, not the programmed values, are listed, followed by an asterisk.

The Usage field represents the current bandwidth assigned to the partition, including guaranteed rate and excess rate for classes with rate policies, and any bandwidth currently allocated to classes with priority policies. Current rate and one-minute averages are bits-per-second rates.

To list only dynamic partitions:

```
PacketShaper# partition show dynamic

Partition name          Size          Grntd       Prior       Curr  1-Min
                        Min /  Max          Excess      Usage       Avg   Peak
-------------------------------------------------------------------------------
 /Inbound/10.2.12.171      0    1.0M     0     0  4608  4608  2804   944  3325
 /Inbound/10.9.50.93       0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.50.47       0    1.0M     0     0     0     0     4    13    53
 /Inbound/10.9.50.27       0    1.0M     0     0     0     0     0     8   214
 /Inbound/152.86.13.0      0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.50.60       0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.50.48       0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.50.44       0    1.0M     0     0     0     0     0     0     0
 /Inbound/8.81.20.0        0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.50.92       0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.51.3        0    1.0M     0     0     0     0     0     9    33
 /Inbound/10.9.50.75       0    1.0M     0     0     0     0     0    39  1943
 /Inbound/114.185.22.0     0    1.0M     0     0     0     0    56    33   302
 /Inbound/10.9.50.1        0    1.0M     0     0     0     0    61    39    62
 /Inbound/10.9.50.109      0    1.0M     0     0     0     0   276   122   655
 /Inbound/10.9.51.13       0    1.0M     0     0     0     0    18    38   238
 /Inbound/114.86.25.0      0    1.0M     0     0  1280  1280   265   178   338
-------------------------------------------------------------------------------
  /Outbound/157/224.0.0.251
                           0    100k     0     0     0     0     0     0     0
```

To list dynamic partitions in a subnet (CIDR=24):

```
PacketShaper# partition show dynamic 10.9.50.0/24

Partition name          Size          Grntd       Prior       Curr  1-Min
                        Min /  Max          Excess      Usage       Avg   Peak
-------------------------------------------------------------------------------
 /Inbound/10.9.50.49       0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.50.93       0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.50.47       0    1.0M     0     0     0     0     0     2    53
 /Inbound/10.9.50.27       0    1.0M     0     0     0     0     0     2   214
 /Inbound/10.9.50.60       0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.50.48       0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.50.44       0    1.0M     0     0     0     0     5     7    27
 /Inbound/10.9.50.92       0    1.0M     0     0     0     0     0     0     0
 /Inbound/10.9.50.75       0    1.0M     0     0     0     0     0    10  1943
 /Inbound/10.9.50.1        0    1.0M     0     0     0     0     9    30   233
```

```
/Inbound/10.9.50.109      0    1.0M     0     0     0     0   535   127   655
```

Command Change History

Release	Modification
8.6.1	**static** and **dynamic** parameters added

ping

Generate pings to test connectivity with another device on the network. If the device answers the pings from the PacketShaper, the message "x.x.x.x is alive" or "x packets transmitted, x packets received" will appear. If PacketShaper is unable to connect with the device, the message "no answer from x.x.x.x" or "0 packets received" will display.

```
ping <host> [<timeout>]
ping [-s] <host> [<count>]
```

<host>	IP address or DNS name
<timeout>	Number of seconds to transmit packets; if you don't specify a *<timeout>* value, PacketShaper will ping the host for up to 10 seconds
[-s]	Send a continuous ping
<count>	Number of pings to transmit; if you don't specify a *<count>* value, PacketShaper will ping the host 10 times

Examples of Successful Pings

PacketShaper# **ping 172.21.1.26**
ping (172.21.1.26): 56 data bytes
172.21.1.26 is alive

PacketShaper# **ping 172.21.1.26 10**
ping (172.21.1.26): 56 data bytes
172.21.1.26 is alive

PacketShaper# **ping -s 172.21.1.26 5**
ping (172.21.1.26): 56 data bytes
64 bytes from 172.21.1.26: icmp_seq=0
64 bytes from 172.21.1.26: icmp_seq=1
64 bytes from 172.21.1.26: icmp_seq=2
64 bytes from 172.21.1.26: icmp_seq=3
64 bytes from 172.21.1.26: icmp_seq=4
5 packets transmitted, 5 packets received

Examples of Unsuccessful Pings

PacketShaper# **ping 192.168.0.1**
ping (192.168.0.1): 56 data bytes
no answer from 192.168.0.1

PacketShaper# **ping 192.168.0.1 30**
ping (192.168.0.1): 56 data bytes
no answer from 192.168.0.1

PacketShaper# **ping -s 192.168.0.1**
ping (192.168.0.1): 56 data bytes
10 packets transmitted, 0 packets received

PacketShaper# **ping -s 192.168.0.1 5**
ping (192.168.0.1): 56 data bytes
5 packets transmitted, 0 packets received

policy admit

Set the admission-control mechanism for a policy.

`policy admit <tclass> squeeze|refuse|"<redirect-URL>" [nontcp|nonweb|web]`

<tclass>	The traffic class whose policy is to be changed
squeeze|refuse| "*<redirect-URL>*"	This admission-control mechanism determines what happens when there isn't enough bandwidth to satisfy a guaranteed rate allocation. When the mechanism is **squeeze**, new connections will get at most 1024 bps. When the mechanism is **refuse**, the connection is refused. For web traffic *only*, when the mechanism is "*<redirect-URL>*", the connection will be redirected to the specified URL.
[nontcp|nonweb|web|all]	The traffic type

The policy admit command supports these combinations:

Admission Control Mechanism	Traffic Types
squeeze	nontcp, nonweb, web
tcp_refuse	nonweb tcp
http_refuse	web
http_redirect	web

policy apply discard

Toss all packets for a class.

```
policy apply discard <tclass>
```

policy apply ignore

Exempt a traffic class from bandwidth allocation and treat the traffic type as "pass-through" traffic.

```
policy apply ignore <tclass>
```

By default, any traffic that you haven't explicitly classified is classified as Inbound/Default or Outbound/Default, and is factored into the bandwidth allocation scheme. When you apply an ignore policy to a traffic class, that traffic type will not be considered at all by the bandwidth allocation process. That is, it won't be counted as part of the virtual link traffic under management.

policy apply never-admit

Force admission control to occur on every use of a policy.

```
policy apply never-admit <tclass>
```

The never-admit policy invokes the appropriate admission-control mechanism at the beginning of each session. For TCP and web traffic, use a never-admit policy to notify users that a service is unavailable. Admission-control mechanisms are configured using the **policy admit** command. For non-TCP traffic, use the **policy apply discard** command. For TCP non-web traffic, you can only use the **policy admit refuse** mechanism with a never-admit policy.

A never-admit policy must be applied to classes on the *requesting* flow. If a never-admit policy is applied to a class representing the response flow, PacketShaper responds as if the policy were a discard policy.

The never-admit policy has proven particularly effective in controlling certain viruses. This type of policy can also be used to redirect certain users to alternate URLs. For example, you might redirect a competitor to a URL that presents a customized message with a competitive analysis. The redirect option works only on the *response side* of the HTTP flow, not the request side.

policy apply priority

Apply a priority-based policy, using the following command:

```
policy apply priority <tclass> [priority]
```

<tclass>	The traffic class to which to apply the policy
[priority]	The priority from 0 to 7, where 7 is highest priority. If this parameter is omitted, priority 3 is used.

Priority-based policies are used to establish a priority for traffic without specifying a particular rate. Use priority policies for traffic that does not burst, or whenever rate is not your primary objective.

policy apply rate

Apply a rate-based policy to a traffic class.

```
policy apply rate <tclass> <guar_lo_bps> <guar_hi_bps> [<priority>[[automatic|<excess_lo_bps><excess_hi_bps>] [<excess_limit_bps>]]]
```

<tclass>	The traffic class to which to apply the policy
<guar_low_bps> *<guar_hi_bps>*	The guaranteed rate for this class' low- and high-speed connections (both parameters are required, even if you specify the same value). Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions). The guaranteed rate must be a minimum of 1024 bps.

For example, to guarantee 10k to Inbound/HTTP, use the following command:

```
policy apply rate inbound/http 10k 10k
```

To allow a policy to use excess rate, specify the following additional parameters:

<priority>	The excess rate priority for this traffic class, ranging from 0 (lowest) to 7 (highest)
[automatic]	Adjusts scaling automatically at run time
[*<excess_lo_bps>* *<excess_hi_bps>*]	The excess rate for this class' low- and high-speed connections (if you don't specify **automatic**). If you choose to use this option, both speeds must be specified. The minimum value allowed for *<excess_lo_bps>* is 1024.
[*<excess_limit_bps>*]	The maximum excess rate that can be used by this class

Guaranteed rate represents the minimum acceptable service level and thus the minimum acceptable rate to allocate. Low- and high-speed rate specifications are used to scale rate allocation to the user's access speed.

For example, to guarantee 10k to Inbound/HTTP burstable up to 48K at priority 3, use the following command:

```
policy apply rate inbound/http 10k 10k 3 automatic 38k
```

> **Note:** Excess rate is expressed differently in the CLI command than in the browser interface. In the browser interface, you specify 48k for the limit, but in the CLI you specify 38k for the amount of excess (the 48k limit minus the guaranteed rate of 10k).

To change the guaranteed rate later, use the policy guaranteed command. To adjust the excess rate, use the policy excess command.

policy default

Apply the PacketShaper-recommended policy to a traffic class.

```
policy default <tclass>
```

policy delaybound

Set the delay bound for a policy to perform non-TCP rate control. PacketShaper uses a *UDP latency control mechanism* to rate-control individual UDP traffic flows and minimize packet loss. PacketShaper accumulates incoming UDP packets on a flow-by-flow basis when they are not scheduled for immediate transfer. With the UDP latency control mechanism, you define a *delay bound* how long the packets can remain buffered before they become too old to be useful. If UDP flows don't get sent immediately (because of link congestion, for example), they are placed in a buffer or *queue*. UDP flows stay in the queue until they are sent or until the delay bound time is exceeded, in which case the packets are dropped.

```
policy delaybound <tclass> [<bound_in_milliseconds>]|default
```

<tclass>	The traffic class whose policy is to be changed
[*<bound_in_milliseconds>*]	The new delay bound, from 1 to 10,000 milliseconds. The default delay bound is set to 200 milliseconds.

> **Note:** Unless you have specific requirements for buffering non-TCP traffic, it is recommended that you do not change the delay bound size, as it has been optimized for most network environments.

Use the **traffic bandwidth** command to view rate exceptions that is, flows that have exceeded the delay bound.

policy dscp

Substitute a value into the Differentiated Services Code Point (DSCP) field in each packet for the class. As defined in the Differentiated Services specification (RFC 2474), the DSCP field is the first six bits of the Type of Service (TOS) field in the IP header. This field is used by routers to make prioritized routing decisions.

```
policy dscp <tclass> unchanged|<dscp>
```

Valid *<dscp>* values are 0-63, inclusive.

policy excess

Modify a rate-based policy's excess rate allocation.

```
policy excess <tclass> <priority> [automatic|<lo_speed_bps> <hi_speed_bps>] [<excess_limit_bps>]
```

<tclass>	The traffic class whose policy is to be changed
<priority>	The new highest priority for excess rate allocation

Optional rate allocations can be specified:

automatic	Automatically scale the low-speed and high-speed rates
[*<lo_speed_bps>*<*hi_speed_bps>*]	The new low- and high-speed rates. If you choose to use this option, both speeds must be specified. The minimum value allowed for *<lo_speed_bps>* is 1024.
[*<excess_limit_bps>*]	The maximum excess rate that can be used by this class

For example, the following command sets the excess rate limits for the FTP traffic class Inbound/Outside/ftp. It is assigned a priority of 4, and assigns both high- and low-speed users an excess rate of 50,000 bps with a total excess rate limit of 200,000 bps:

```
policy excess /inbound/outside/ftp 4 50k 50k 200k
```

policy flowlimit

Limit the rate of new flows to or from a unique host. This command can be used to detect and control a SYN Flood or similar denial-of-service attack directed at a particular host or if the attack is from a specific IP address. Flows exceeding the rate are blocked from passing through the unit. The limits are set to default values of 10,000 flows per minute on client hosts (flow initiators) and 100,000 flows per minute on servers (flow destinations); depending on your network, you may need to change these defaults for effective control of SYN floods. Flow limits are automatically set on any classes that have a rate or priority policy assigned to them; if the **PolicyFlowLimitForAllClasses** system variable is enabled, PacketShaper will automatically

block any flows that exceed these limits. (This variable is disabled by default. If you want to enforce flow limit policies, you need to enable the **PolicyFlowLimitForAllClasses** variable.)

> **Note:** You cannot set a flow limit on a class unless it already has a rate or priority policy assigned to it.

If you want to set or adjust the default limits on a particular class, use:

```
policy flowlimit <tclass> none|<client-fpm> <server-frm>
```

<tclass>	Traffic class where the policy is located
none	Remove the flow limit
<client-fpm>	Maximum number of flows per minute to allow from each individual host (the flow initiator); valid values are 0-600000
<server-fpm>	Maximum number of flows per minute to allow to each individual host (the flow destination); valid values are 0-600000

Note that the *<client-fpm>* and *<server-fpm>* rates include new flows of all types from an individual client (flow initiator) or to an individual server (flow destination) —not just flows of the type of traffic matching this specific traffic class or policy. For example, *<client-fpm >* includes flows from a particular client to many point destinations. And *<server-fpm>* includes flows from many point sources to a particular server.

PacketShaper offers measurement variables to track the number of flows that were blocked due to a server (flow destination) or a client (flow initiator) exceeding the flow limit rate specified in the **policy flowlimit** command: server-flood-block and client-flood-block.

If you don't want flow limits to be set automatically for newly created classes, enter the following commands:

```
policy flowlimit inbound/default none
policy flowlimit outbound/default none
```

policy guaranteed

Modify a rate policy's guaranteed rate allocations.

```
policy guaranteed <tclass> <lo_speed_bps> <hi_speed_bps>
```

<tclass>	Traffic class whose policy is to be changed
<lo_speed_bps> *<hi_speed_bps>*	New low-speed and high-speed guaranteed rates. Rates may be specified as integer bits per second, followed by a k (thousands), M (millions), or G (billions). The guaranteed rate must be a minimum of 1024 bps.

For example, the following command sets the low-speed and high-speed rates (10000 bps for low-speed users and 100000 bps for high-speed users) for a class named inbound/jup_202_http:

```
policy guaranteed inbound/jup_202_http 10000 100000
```

policy precedence

Substitute a precedence value for IP-based traffic classes.

```
policy precedence <tclass> unchanged|<precedence>
```

<tclass>	The traffic class for which you want to change precedence.
unchanged | *<precedence>*	Use **unchanged** to turn off precedence substitution, restoring precedence to its default value. Or, enter a precedence value 0-7, where 7 is the highest priority.

> **Note:** The **policy precedence** command supplements rate and priority policies that is, a traffic class must have a policy already applied to it before you use the **policy precedence** command to substitute a precedence value.

policy remove

Remove a policy from a traffic class.

```
policy remove <tclass>
```

policy route

Divert specific traffic to an alternative route by sending the class' traffic to a secondary gateway or router.

Set the MAC address routing for a traffic class.

```
policy route <tclass> none|<macaddr>
```

PacketShaper substitutes the MAC address and transmits the packet accordingly.

policy show

Display policy information.

```
policy show <tclass> [clear]
```

<tclass>	Explicit traffic class name whose policy is to be displayed - for example, Inbound/Outside/http
[clear]	Reset the associated traffic class and policy hit counts

policy substitute

PacketShaper can detect the speed of a web connection at the first HTTP get request. You can use the **policy substitute** command to re-map the requested URL by substituting a URL that's more appropriate for the speed of the connection.

```
policy substitute <tclass> none
```

```
policy substitute <tclass> above|below <speed> "<pattern>" "<newpattern>"
```

<tclass>	The traffic class to which you are applying the policy
above|below	Specify **above** or **below** a connection speed to indicate when the URL should be substituted.
<speed>	The connection speed that, in conjunction with **above** or **below**, triggers the content substitution
"<pattern>"	Specify in quotes the current URL pattern, which will be substituted with a new pattern. Wildcard patterns are not supported. This URL string is compared with the pattern in the /directory/file portion of a URL. PacketShaper ignores the http://computer-name portion of a URL when performing matching or substitution.
"<newpattern>"	The URL that you specify for substitution *must be the same length as the original URL*. The formatting rules are the same as those listed for the *<pattern>* parameter.

For example, to better serve a low-speed user, you could substitute a text-based web page for the regular home page:

```
policy substitute inbound/outside/web-in below 28800 "home-1.htm" "home-2.htm"
```

policy test

Test a policy to determine what rate will be allocated.

```
policy test <tclass> <rate_bps>
```

<tclass>	The traffic class whose policy is to be tested
<rate_bps>	The access speed to use to determine rate allocation

Example:

Assume the class inbound/http has the following policy settings: 10k guaranteed, burstable at priority 5, limit of 100k. To see how excess rate is allocated when there is 150 Kbps of demand, use the following command:

```
policy test inbound/http 150k
Policy Settings
 Guaranteed rate lo 10k      hi 10k
 Excess rate     default priority 5  CAP 90k
Allocation for flow at rate 150000
 Guaranteed rate 10000
 Excess rate at priority 5 -> 25088
 Excess rate total 90000
 Excess rate demand 0       0       15k     25k     25k     25k     0       0
```

This output shows how PacketShaper would allocate bandwidth when traffic class inbound/http generates 150 Kbps of demand. The top part of the display summarizes the policy settings. The excess rate (90k, next to CAP) is calculated by subtracting the guaranteed rate from the limit (100k-10k=90k).

The lower portion of the output lets you see how the excess rate is allocated between priority levels, 0-7. The sum of the rates allocated at each priority level equals the total excess rate (90k, in this example).

policy tos

Set a specific type of service for an IP traffic flow. It can be applied only to a class that already has a rate control policy defined.

`policy tos <tclass> unchanged|<tos>`

<tclass>	Explicit traffic class name for which you want to change the type of service
unchanged |*<tos>*	Use **unchanged** to turn off TOS substitution. Enter a *<tos>* value according to the following standard: 8 = minimize delay 4 = maximize throughput 2 = maximize reliability 1 = minimize monetary cost 0 = normal service Values can be combined to define broader results. For example, a value of 3 indicates "maximize reliability and minimize monetary cost."

policy vlan

Add or change a VLAN identification (802.1Q) or priority (802.1p) on a packet. It can be applied only to a class that already has a rate or priority policy defined.

VLAN Priority (802.1p)

To change the priority tag on an 802.1p class:

`policy vlan type:8021p <tclass> swap <priority>`

<tclass>	Traffic class for which you want to modify the VLAN priority tag
swap *<priority>*	Swaps the topmost priority level on the VLAN stack with the specified *<priority>*, 0 to 7.

For example, to change the VLAN priority to 6:

```
policy vlan type:8021p vlantestclass swap 6
```

VLAN Identification (802.1Q)

To modify the identification tag on an 802.1Q class:

```
policy vlan type:8021q <tclass> push|swap <vlanid> | pop <times> | delete <operation>
```

<tclass>	Traffic class for which you want to modify the VLAN ID
push *<vlanid>*	Puts an ID entry in a packet (and creates the stack if it doesn t exist); the *<vlanid>* is the value of the label to be pushed (0 4095)
swap *<vlanid>*	Swaps the topmost ID of the VLAN stack with the specified *<vlanid>* (0 4095)
pop *<times>*	Pops off the topmost label of the VLAN stack in the packet the specified number of *<times>*
delete *<operation>*	Deletes the specified *<operation>* (pop, swap, or push) from the policy

Examples:

```
policy vlan type:8021q testclass pop 2

policy vlan type:8021q testclass push 1

policy vlan type:8021q testclass swap 6

policy vlan type:8021q testclass delete pop
```

A class can have a combination of push, swap, and pop operations in its VLAN policy; the pop operation can be specified multiple times (up to 8). If more than one operation type is specified for a given class, they are executed in the following order: pop, swap, push. For example, a class might have a policy that specifies a swap, three pops, and a push. In this case, the three pops occur first, then the swap, and then the push.

> **Note:** A VLAN ID swap policy will automatically zero out the existing VLAN priority. To keep an existing non-zero priority value or to set a priority, be sure to specify a VLAN priority swap policy as well.

portal delete

Delete a customer portal account.

```
portal delete <name>|all
```

Command Change History

Release	Modification
11.6.1	Command introduced

portal modify

Modify customer account information.

```
portal modify <name> <password> <directory> <message-of-the-day>
```

If RADIUS or TACACS+ authentication is enabled, passwords are not used (they are entered at the RADIUS/TACACS+ server). Thus, the syntax when RADIUS or TACACS+ is enabled is:

```
portal modify <name> <directory> <message-of-the-day>
```

Parameter	Description
<name>	The existing customer login name
<password>	The password for the customer account. If you are using RADIUS or TACACS+ authentication, you do not specify a password here—the customer portal will use the password specified for this user in the RADIUS/TACACS+ server.
<directory>	The new name of the customer's home directory (up to 8 characters); this directory will be created on the unit's drive under 9.258/customer
<message-of-the-day>	The new custom message-of-the-day (optional)

If you don't specify the parameters, PacketShaper will prompt you for the information:

```
portal modify
```

Enter the name of the customer : **mycust**
Enter the new password :
Confirm new password :
Enter the new home directory name, 8 characters or less : **newdir**
Enter the new custom message-of-the-day (optional) : **All network resources online**
Customer mycust was modified

> **Note:** You will not be prompted for a password if RADIUS authentication is enabled.

After this is executed, mycust's home directory will be 9.258/customer/newdir.

> **Note:** You must explicitly type each entry when you use prompted mode. If, for example, you press Enter at the password prompt, the new password value becomes (none).

Command Change History

Release	Modification
11.6.1	Command introduced

portal new

Create a new customer portal account.

```
portal new <name> <password> <directory> <message-of-the-day>
```

If RADIUS or TACACS+ authentication is enabled, passwords are not used (they are entered at the RADIUS/TACACS+ server). Thus, the syntax when RADIUS or TACACS+ is enabled is:

```
portal new <name> <directory> <message-of-the-day>
```

Parameter	Description
<name>	The login name the customer will use; up to 32 characters long, use numbers, letters and underscores spaces are not allowed. If you are using RADIUS or TACACS+ authentication, this name must match the user name entered in the RADIUS/TACACS+ server.
<password>	The password for the customer account. If you are using RADIUS or TACACS+ authentication, you do not specify a password here– the customer portal will use the password specified for this user in the RADIUS/TACACS+ server.
<directory>	The name of the customer's home directory (up to 8 characters); this directory will be created on the unit's drive under 9.258/customer
<message-of-the-day>	A text string of 128 characters or less, intended to carry simple messages such as *System will be down from 5:00 am to 6:00 am* tomorrow (optional)

You must use empty quotes ("") if you don't want to enter a value for a parameter. For example, to create a user MyCust with a directory named cust01 (no password, no message of the day), use:

```
portal new MyCust "" cust01 ""
```

If you don't specify any parameters with the **portal new** command, PacketShaper will prompt you for the values.

This is an example of prompted mode:

```
portal new
```

Enter the customer login name, password, home directory name (8 characters or less) and an optional custom message-of-the-day (128 characters or less).
Enter the customer's login name, e.g. 'marysmith' : **mycust**
Enter the password :
Confirm the password :
Enter the customer's home directory name, e.g. 8 characters or less : **mycust**
Enter a custom message-of-the-day (optional): **No network outages**
Customer mycust was added.

> **Note:** You will not be prompted for a password if RADIUS authentication is enabled.

After this is executed, a directory 9.258/customer/mycust will exist. The service provider must FTP an INDEX.HTM file to it before the mycust customer can use it effectively.

> **Note:** You must explicitly type each entry when you use prompted mode. If, for example, you press Enter at the password prompt, the new password value becomes (none).

Command Change History

Release	Modification
11.6.1	Command introduced

portal show

Display the current customer portal configuration.

```
portal show
```

Customer Name	Password	Directory	Message
Farnsworths	******	books	Inventory starts Friday!
Sigma_Air	******	air	No scheduled network outages

The Password column does not appear if RADIUS authentication is enabled.

Command Change History

Release	Modification
11.6.1	Command introduced

pwd

Show the working directory.

```
pwd
```

radius acct

Test and debug the setup of your RADIUS accounting server. This command sends test accounting messages to the server.

```
radius acct start|stop|on|off
```

Specify start to send a test message that tells the accounting server that someone logged in and stop to send a log-off test message. The administrator can then verify that these messages are in the accounting server log. They will appear in the log under the name *RadiusAccountingTestUser*.

You can use the **on** and **off** parameters to send a message to the server that the RADIUS accounting service is on or off. Note that this command does not affect the setup of the accounting service; if the service was enabled with the **setup radius acct on** command, it will remain enabled (even if you used the **radius acct off** to send a test message that the accounting service is off).

Command Change History

Release	Modification
11.5.1	Command introduced

radius chaplogin

Send a test CHAP login request to the RADIUS authentication server. This command is useful for testing and debugging the setup of your RADIUS authentication server, when Challenge-Handshake Authentication Protocol (CHAP) is used.

```
radius chaplogin <username> <password>
```

For example:

```
radius chaplogin bob 12567
```

```
chap ID = 0x1
 challenge = 37a9aa04189c7ac5c826fde6a52c988f
 password = 12567
 response = 7610c93540dc90422fb4b077d23dd63a
 "bob" RADIUS Authentication OK
 Vendor-Specific: access=touch
```

The above output indicates that the authentication of the user Bob was successful. If authentication fails, you will see one of the following messages:

Message	What it means	What you should do
Authentication turned off	You need to enable the authentication service on PacketShaper	Use the **setup radius auth on** command to enable authentication, and then send another test login request.
No server configured	The RADIUS authentication service is turned on in PacketShaper, but the server is not configured.	Use the **setup radius auth primary** command and specify the authentication server's IP address, port number, and shared secret. Then send another test login request.
Access rejected by server	The user name and/or password is invalid.	Contact your RADIUS administrator to verify that you are using the correct user name and password.

Message	What it means	What you should do
Timeout: Unable to obtain a response from server	The RADIUS authentication service is turned on in PacketShaper and the server is configured. This message could be caused by any of the following situations:	
	Server could be down.	Contact your RADIUS administrator to check the status of the RADIUS authentication server. It's a good idea to configure a secondary server to have a backup in case the primary server fails.
	Incorrect IP address for the server.	Contact your RADIUS administrator to verify the host name or IP address of the authentication server.
	The authentication service may not be enabled on the RADIUS server side.	Contact your RADIUS administrator to verify that the authentication service is enabled on the RADIUS server.
	The server may not be configured to work as a PacketShaper client.	For information on configuring the RADIUS server with PacketShaper-specific attributes, see Configure RADIUS Servers in PacketGuide.
	The LAN may be busy or down.	Check the status of the network.

Command Change History

Release	Modification
11.5.1	Command introduced

radius clear

Clear the accounting drop count and remove the drop-notice banner. When an accounting request is dropped because the accounting server was not configured correctly or was unreachable for some reason, PacketShaper keeps track of these dropped accounting requests and displays a banner alerting you that requests have been dropped. You can use the **radius clear** command to clear this banner.

```
radius clear
```

Command Change History

Release	Modification
11.5.1	Command introduced

radius login

Send a test PAP login request to the RADIUS authentication server. This command is useful for testing and debugging the setup of your RADIUS authentication server, when Password Authentication Protocol (PAP) is used.

```
radius login <username> <password>
```

For example:

```
radius login bob 12567
```

"bob" RADIUS Authentication OK
Vendor-Specific: access=touch

The above output indicates that the authentication of the user Bob was successful. If authentication fails, you will see one of the following messages:

Message	What it means	What you should do
Authentication turned off	You need to enable the authentication service on PacketShaper	Use the **setup radius auth on** command to enable authentication, and then send another test login request.
No server configured	The RADIUS authentication service is turned on in PacketShaper, but the server is not configured.	Use the **setup radius auth primary** command and specify the authentication server s IP address, port number, and shared secret. Then send another test login request.
Access rejected by server	The user name and/or password is invalid.	Contact your RADIUS administrator to verify that you are using the correct user name and password.
Timeout: Unable to obtain a response from server	The RADIUS authentication service is turned on in PacketShaper and the server is configured. This message could be caused by any of the following situations:	
	Server could be down.	Contact your RADIUS administrator to check the status of the RADIUS authentication server.
		It s a good idea to configure a secondary server to have a backup in case the primary server fails.
	Incorrect IP address for the server.	Contact your RADIUS administrator to verify the host name or IP address of the authentication server.
	The authentication service may not be enabled on the RADIUS server side.	Contact your RADIUS administrator to verify that the authentication service is enabled on the RADIUS server.
	The server may not be configured to work as a PacketShaper client.	For information on configuring the RADIUS server with PacketShaper-specific attributes, see Configure RADIUS Servers in PacketGuide.
	The LAN may be busy or down.	Check the status of the network.
Error: Reply didn't contain an access level attribute	The user name and password are valid, but the user wasn't configured with an access level attribute.	Configure the RADIUS server with an access level attribute for this user.

Command Change History

Release	Modification
11.5.1	Command introduced

radius session

Show a list of current user sessions (RADIUS, TACACS, ds, local) and detailed information about each session.

```
radius session
```

ID	Status	Age	Idle	Limit	Type	Access	User Name
3b61a571	logged in	40 mins	0 secs	60 mins	CLI	touch	bob
3b61afb6	logged in	91 secs	91 secs	60 mins	WUI	touch	john
3b61af6a	logged in	167 secs	17 secs	60 mins	WUI	touch	george
3b61a571	logged in	45 mins	0 secs	60 mins	CLI	touch	bob

Column	Description
ID	Identification given to the user session
Stat	The status of the session: logged in– the session is active timed out– the session timed out logged off– the session is inactive
Age	Length of time the session has been active– that is, the amount of time since the user logged in
Idle	Amount of time since the user gave a command; whenever a user gives a command, the idle value is reset to zero
Limit	Amount of time a session is idle before the user will be timed out and logged off; for example, if the limit is 60 minutes, a user will get logged off when no commands are given for a 60-minute period. **Note:** The PacketShaper default session life limit is 60 minutes. However, the RADIUS or TACACS+ server can be independently configured with different limits for different users and these limits override PacketShaper s.
Type	Type of interface used: CLI (command-line interface), WUI (web user interface)
Access	Type of access: Touch, Look
User Name	Name of the user who logged into the session

Command Change History

Release	Modification
11.5.1	Command introduced

radius show

Display RADIUS client configuration. Use this command to verify that RADIUS authentication and accounting are enabled, to see the current settings for the retry limit and retry interval, and to view the configuration settings for the primary and secondary authentication and accounting servers.

```
radius show

Radius method is CHAP
Radius Authentication is ON
Radius Accounting is OFF
Retry limit: 3
Retry interval: 5

         auth1            auth2            acct1            acct2
-------------------------------------------------------------------------------
Server   172.23.225.203   172.23.225.213   -                -
Secret   packet           packet           -                -
Port     1812             1812             0                0
Status   Up               Unknown          Unknown          Unknown
Attempts 1                0                0                0
Success  1                0                0                0
Timeout  0                0                0                0

Auth1 last accessed: Wed Jul 11 14:16:48 2007
Auth2 was never accessed!
```

The output also indicates the number of attempts made to connect to each server, the number of successful connections, and the number of connections that timed out.

Command Change History

Release	Modification
11.5.1	Command introduced

reset

Reset the unit. In certain situations, you may need to reboot the unit, for example, after creating host accounting categories and changing system variables . A reset performs a proper UNIX reboot, without losing power to the sytstem.

```
reset
```

The statistics that are cleared after you reset the device include: active IP hosts, current speed per host, active flows, peak flows, and Top Talkers/Listeners. All counters on the Monitor Traffic page are reset. All traffic configurations (such as classes, policies, and partitions) are preserved when you reboot.

To reset all settings to factory default, use the setup reset all command.

rtm accept

Set an acceptable service level threshold percentage for response time measurement (RTM). The default is 100%.

```
rtm accept <tclass> <percent>
```

where *<tclass>* is the traffic class to be defined.

Command Change History

Release	Modification
11.4.1	Command introduced

rtm clear

Zero out response time measurement statistics for all classes.

```
rtm clear
```

Command Change History

Release	Modification
11.4.1	Command introduced

rtm drilldown

List the hosts with the highest percentage of slow transactions, as defined by the class' Total Delay Threshold. A feature of the drilldown command is to suppress hosts that had fewer than N transactions, as they might skew the data.

```
rtm drilldown <tclass> [<number> [<cutoff>]]
```

<tclass>	Specify a traffic class that tracks response time and has a threshold.
<number>	Limit the number of displayed entries (default is 10).
<cutoff>	Don't list hosts with fewer than *cutoff* transactions. This eliminates the hosts whose response figures aren't meaningful because too few data points were available (few transactions). If you want to include every host in the **rtm drilldown** output, set the optional cutoff parameter to 1.

Command Change History

Release	Modification
11.4.1	Command introduced

rtm hosts

Enable or disable worst client/server tracking for a class.

```
rtm hosts <tclass> enable|disable
```

where *<tclass>* is the traffic class to be tracked.

Command Change History

Release	Modification
11.4.1	Command introduced

rtm show

Display a summary of the RTM statistics for all traffic classes with response-time data.

```
rtm show
```

The display has one row per traffic class with the following information:

Traffic Class	The name of the traffic class
Goodness	The number of good transactions (those within the Total Delay Threshold) divided by the transaction count, multiplied by 100. In other words, the percentage of good transactions.
Response Time: Total, Network, Server, Normal	The average number of milliseconds required by the class' transactions. The value in the Normal column is the component of the transaction time that is directly related to the transaction size. An increase or decrease in that number does not indicate any change in network or server performance and requires no user intervention. This value is not tracked by the measurement engine.

Command Change History

Release	Modification
11.4.1	Command introduced

rtm threshold

Differentiate between acceptable and unacceptable response by supplying a threshold that defines good performance. PacketShaper uses the threshold when evaluating each transaction's total delay figure. If the transaction completes within the time indicated with the threshold, the transaction is considered "good."

```
rtm threshold <tclass> <delay>|none [total|network|server]
```

Specify the delay threshold in milliseconds or remove the threshold using the **none** literal. The threshold maximum is 99 seconds.

If you set a network or server delay threshold in the CLI, the setting will not appear in the browser interface, as these types of thresholds are not supported in the browser version of PacketShaper. (Only total delay threshold can be set in the browser interface.) Note that using the browser interface to make any changes to a class RTM settings will clear the network or server delay you set in the CLI.

Command Change History

Release	Modification
11.4.1	Command introduced

rtm worst

Display the traffic classes that have the "worst" response-time statistics.

```
rtm worst [goodness|total|network|server [<number> [<cutoff>]]]
```

goodness|total|network|server	Displays response time measurement data sorted by Total Delay, Network Delay, Server Delay, or Goodness. If a data type is not specified, goodness is the default type.
<number>	Limits the number of classes that are displayed. The default value is 10.
<cutoff>	Excludes the traffic classes that have less than the number of transactions specified by this cutoff value. The default is 10.

Command Change History

Release	Modification
11.4.1	Command introduced

run

Run a command file or script. This command runs in the context of your current folder. The filename must have a **cmd** suffix. An output file, *<filename>*.out, contains the results of command-file execution. To view the contents of this output file, use the more command.

> **Note:** PacketShaper's diagnostic commands (dns, ping, net, sys, uptime) and utility commands (cd, cmp, date, history, ls, more, pwd) cannot be executed from a command file or used with the **schedule** command.

schedule delete

Delete a scheduled command execution.

```
schedule delete <item_id>|all
```

The scheduled item is removed from the list, but the remaining items have the same IDs.

This command produces output similar to the following:

Id	Time Range	Issued	Date	Command
00000001:	08:00:00-08:00:40	1	weekday	"p2pday.cmd"*
2DA2B03E:	18:00:00-18:00:40	1	weekday	"p2peve.cmd"*
5545F94E:	06:00:00-06:00:40	0	daily	"test.cmd"*

schedule disable

Disable a scheduled item so that it won't be executed. If you want to permanently remove the scheduled item, use the schedule delete command.

```
schedule disable <item_id> | all
```

where *<item_id>* is the ID of the item displayed in the schedule show output. Use the **all** parameter to disable all entries.

After you have disabled an item, it will still be listed in the **schedule show** output, but [disabled] will appear at the end of the line.

Id	Time Range	Issued	Date	Command
00000001:	08:00:00-08:00:40	1	weekday	"p2pday.cmd"*
2DA2B03E:	18:00:00-18:00:40	1	weekday	"p2peve.cmd"*
7EC07402:	06:00:00-06:00:40	1	weekend	"p2psat.cmd"*
5545F94E:	06:00:00-06:00:40	0	daily	"test.cmd"*

To enable the item later, use the schedule enable command.

schedule enable

Enable a scheduled item that has been disabled with the schedule disable command.

```
schedule enable <item_id> | all
```

where *<item_id>* is the ID of the item displayed in the schedule show output. Use the **all** parameter to enable all entries.

schedule new

Schedule a command to execute at a specific time and date. When using the scheduling feature, it s important that your unit has the correct date and time. Use the date command to check the date and time. If you need to correct the date or time, use the setup date command.

Scheduling is limited to 64 scheduled commands. Scheduled commands that are no longer needed (for example, expired commands scheduled to run only once) should be removed from the list via the **schedule delete**command, so they do not continue to consume available resources.

Note for units in shared mode: To permanently delete a command scheduled via PolicyCenter, you must remove the command from the unit's sharable PolicyCenter configuration, and not just from the unit itself. If PolicyCenter detects that a scheduled command is on the unit's PolicyCenter configuration but no longer on the unit itself, PolicyCenter will synchronize the unit's settings with its PolicyCenter configuration to restore that command to the unit. If the time range for the scheduled command is not over or the PolicyCenter time zone is not correctly configured, the command may run again.

When setting the time for a scheduled command, keep biannual time changes into consideration. For example, if you set a command to execute at 2:30am in the United States, the command will not be executed when the clock changes ahead one hour for Daylight Saving Time. You can ensure a command will be executed during a clock change by specifying a time range (such as 02:00-04:00).

```
schedule new [<day option>] <time_range>[utc] <cmd>|{-f <cmd_file>} [mail:<address>]
[id:<item_id>] [disable]
```

[*<day option>*]

Specifies the day(s) the schedule should run. If you don t specify the *<day option>*, the scheduled item will run every day.

Specify one of the following for *<day option>*:

now

today | + *<n>*

[once:]*<date>*[,*<date>*]

[once:]weekday | weekend | *<dow>*[,*<dow>*]

[once:]*<dom>*[,*<dom>*]

+*<n>* is *<n>* days from today. For example, **+1** is tomorrow.

<date> is a specific date in the format *mm/dd*. The date is assumed to be a future date, within the next twelve months. For example, if today is 5/30/02 and you specify the date 5/29, the item will be scheduled for execution on May 29, 2003. You can specify up to 10 dates separated by commas.

The **once:** option, that optionally precedes the *<date>* and the following options, specifies that the item should be executed once for each of the specified dates. If you don t specify **once:**, the item will be executed on an ongoing basis, according to the date(s) you specified.

weekday executes the item on weekdays only (Monday through Friday). **weekend** executes the item on weekends (Saturday and Sunday). These two options are useful for setting different policies for weekdays and weekends. For example, you might want music file sharing to have less bandwidth during the week than on the weekend.

<dow> is the day of the week, specified with the first three letters of the day (mon, tue, wed, thu, fri, sat, sun). If you specify more than one day, each day is separated by a comma and no space, for example, *mon,wed*. You can specify up to seven days of the week.

<dom> is a specific day of the month, for example, *15* for the fifteenth of the month. You can specify up to 31 days, separated by commas.

<time_range> [utc]

Specifies the time at which the command or command file should be executed. The syntax is:

```
hh:mm[.ss][-hh:mm[.ss]
```

where *hh* is the hour from 0 to 23, and *mm* and *ss* are minutes and seconds from 0 to 59. For example, to specify the time 5pm, enter *17:00*.

If a range is not specified, PacketShaper will attempt to execute the command within a 40-second window. If you want to allow more time for the command(s) to be executed, you can specify a range (for example, 08:00-08:02).

The legacy syntax *hhmm.ss* is supported for backward compatibility.

Use the optional suffix **utc** to specify a coordinated universal time (UTC). Specifying times in UTC (similar to Greenwich mean time) is useful when managing units in different time zones. For example, 1800Z is 1pm in Eastern standard time and 4pm in Pacific standard time.

<cmd> | -f *<cmd_file>*	You can specify one of the following: ■ The CLI command *<cmd>* to be executed. The command should be enclosed in quotation marks. ■ The name of the file (*-f <cmd_file>*) that contains a list of CLI commands. Specify a path to the *<cmd_file>* unless the file is in the default folder (9.256/cmd). The filename should be eight characters or less and have a .CMD extension.
[mail: *<address>*]	Sends the output of the command or command file execution to the specified email address(es), allowing you to confirm that the command executed at the specified date and time. You can specify up to four email addresses, separated by commas. In order to use this feature, you must configure a mail server. See setup email.
[id: *<item_id>*]	Assigns the specified ID to the scheduled item. *<item_id>* can be up to eight characters long and can contain the numbers 0-9 and the letters A-F and a-f. This parameter is primarily used to override inherited entries when using shared mode (PolicyCenter). The ID is shown in the list of scheduled entries via the schedule show command. If you don't specify an ID, PacketShaper assigns a random number. **Note:** It's recommended that you allow PacketShaper to automatically create the ID rather than manually assign the ID with the **id** option. If you do manually assign an ID, make sure you follow the guidelines for ID names, as described above.
[disable]	Disables the scheduled item so that it won't be executed. If you want to enable or disable the item after it is created, use the schedule enable or schedule disable command.

You will typically want to create scheduled items in complementary pairs. For example, you can create one scheduled item for a policy that is applicable during work hours and another schedule for a policy that is applicable after hours.

```
schedule new weekday 08:00 policy apply rate /inbound/gnutella 4800 9600 2 4800 4800
 schedule new weekday 18:00 policy apply rate /inbound/gnutella 128k 256k 4 256k 256k
```

If you use the mail:*<address>* parameter, an email message containing the command output will be sent to the specified *<address>* shortly after the schedule is executed.

> **Note:** PacketShaper's diagnostic commands (dns, ping, net, sys, uptime) and utility commands (cd, cmp, date, history, ls, more, pwd) cannot be executed from a command file or used with the **schedule** command.

Command Change History

Release	Modification
11.5.1	Support for [mail:*<address>*] option added
11.1.1	Support for [mail:*<address>*] option removed
8.3.1	**[d]** parameter to automatically delete command file removed
8.0.1	**now***<day option>* introduced
8.0.0	no change

schedule show

List the currently scheduled commands.

```
schedule show [-time] [-utc]
```

[-time] Sorts schedules by time, with the earliest start time listed first.
Without the **-time** switch, schedules are sorted by ID.

[-utc] Lists schedules in their original time input. If time was entered in coordinated universal time (UTC) format, the UTC time will be displayed with a Z after the time (for example, 00:00:00-00:00:40Z). If the time was specified in local time, the local time will be displayed.
Without the **-utc** switch, all times are listed in local time; in other words, any UTC times are converted to local time on the display. An L displays after the converted UTC time (for example, 08:00:00-08:00:40L).

This command produces output similar to the following:

Id	Time Range	Issued	Date	Command
00000001:	08:00:00-08:00:40	1	weekday	"p2pday.cmd"*
2DA2B03E:	18:00:00-18:00:40	1	weekday	"p2peve.cmd"*
7EC07402:	06:00:00-06:00:40	1	weekend	"p2psat.cmd"*
5545F94E:	06:00:00-06:00:40	0	daily	"test.cmd"*

Each scheduled event has a unique ID, which can be used to delete items from a schedule. The number in the "Issued" column indicates how many times the command has executed. An asterisk (*) flags command-file items. If [disabled] appears, the item has been disabled with the schedule disable command and will not be executed.

send email

Define an email message. Include this command in an adaptive response action file to trigger an email notification of a change in the status of a unit or your network.

```
send email <address> "<subject>" ["<body>"]
```

<address> Email address of the recipient; up to two addresses can be specified, separated by a space

<subject> Text to be included on the subject line

<body> Text to be included in the body of the email message

The total maximum length of the **send email** command is 256 characters.

> **Note:** You must first configure a mail server with the **setup email** command before you can issue the **send email** command.

Command Change History

Release	Modification
11.5.1	Command introduced

send syslog

Define a syslog message that will be sent to a syslog server.

```
send syslog FDBK <severity><mnemonic> ["<string>"]
```

<severity>	Specify one of the following severity levels: Emergency, alert, critical, error, warn, notice, info, debug See Severity Levels for more information about security levels.
<mnemonic>	Any 9-character string that can be used to categorize the message
"*<string>*"	Text to be included in the syslog entry

> **Note:** You must configure a syslog server with the setup syslog command before you can use you can use **send syslog** command to send an syslog message.

Including this command in an adaptive response action file triggers a syslog message to be sent to a previously defined syslog server. If you include action file variables in this command, the agent automatically enters the values for the variables as the action file is run. The user cannot change the values of these variables; their values come from the agent only.

The following example command could be included in an action file for a "High Bandwidth Host" agent:

```
send syslog FDBK error $agentname "$scorecolor: $namelist is the biggest violator."
```

The High Bandwidth Host agent will automatically fill in the *$agentname*, *$scorecolor*, and *$namelist* variables with the name of the agent, the agent's status color at the time the action file was triggered, and the name of the host.

send trap

Define an SNMP (Simple Network Management Protocol) trap that will be sent to an SNMP trap listener.

```
send trap "<name>" <color> ["<description>"]
```

<name>	Name of an agent for which you want to send a trap. If the agent name has a space, enclose it in quotation marks, for example, *"My Agent."*
<color>	Agent's status color at the time of the trap: red, yellow, green, or blue
<description>	Description of what happened at the time of a trap

Including this command in an adaptive response action file triggers a trap to be sent to a previously defined SNMP trap destination. When you include adaptive response agent *action file variables* in this command, the agent automatically enters

the values for the variables as the file is run. The user cannot change the values of these variables; their values come from the agent only.

The following example command could be included in a red action file for a Packet Drops agent:

send trap "$agentname" $scorecolor "$classname is the biggest violator."

The Packet Drops agent will automatically fill in the *$agentname*, *$scorecolor* and *$classname* variables with the name of the agent, the agent's status color at the time of the trap, and the name of the traffic class. If excessive packet drops in the traffic class /Inbound/HTTP triggered the red action file, the action file variables would made the action file read as follows:

send trap "Packet Drops" red "/Inbound/HTTP is the biggest violator"

> **Note:** To send SNMP traps, PacketShaper needs to know where to send the traps.

setup access default

When issued from the command-line interface of an individual unit, this command returns the security access settings for the unit to its default value, allowing users to access the unit via all available secure and non-secure protocols.

> **Note:** To enable or disable a single protocol for accessing the unit, use instead the commands setup access enable or setup access disable.

setup access default

setup access disable

Allow access to a PacketShaper's browser and command-line interfaces (CLI) only via specific access protocols or with a secure connection. By default, the secure access protocols (HTTPS, SSH) are enabled and the non-secure protocols (HTTP, SNMP) are disabled.

> **Note:** Changing this setting might cause the session to be dropped.

setup access disable all|[https|ssh|http|snmp]

where:

all	Disable all secure and nonsecure protocols for accessing your unit. Note that if you disable all protocols, you will only be able to access the unit via a direct console connection.
https	HTTP over Secure Sockets Layer protocol (SSL). The browser interface uses the SSL protocol to securely access the unit. Disabling this protocol will disable secure access to the unit via the browser interface.
ssh	Secure Shell remote login protocol (SSH). The CLI uses SSH to securely access the unit. Disabling this protocol will disable secure access to the unit via the CLI.
http	Hypertext Transport Protocol. This is a non-secure protocol for the browser interface.
snmp	Simple Network Management Protocol. This is a non-secure protocol.

setup access enable

Issue this command to reenable access to the unit via a service protocol that was previously disabled with the setup access disable command. By default, all services are enabled, allowing you to access the unit by all available secure and non-secure protocols.

setup access enable all|[https|ssh|http|snmp]

Where:

all	Enable all secure and nonsecure protocols for accessing your unit.
https	HTTP over Secure Sockets Layer protocol (SSL). The browser interface uses the SSL protocol to securely access the unit.
ssh	Secure Shell remote login protocol (SSH). The CLI uses SSH to securely access the unit.
http	Hypertext Transport Protocol. This is a non-secure protocol.
snmp	Simple Network Time Protocol. This is a non-secure protocol.

Command Change History

Release	Modification
11.5.1	Added SNMP access
11.1.1	Eliminated support for ftp, telnet, echo, and snmp access

setup access show

Display access security settings for the unit. The output lists all service protocols available for accessing the unit, and indicates whether each protocol is enabled or disabled.

setup access show

To disable or reenable a protocol for accessing the unit, use the command setup access disable or setup access enable.

For example:

```
Unit Access:
Access Method    Disabled
-----------------------------
HTTPS            No
SSH              No
HTTP             Yes
SNMP             No
```

setup adaptiveresponse

Turns *all* configured adaptive response agents on or off, or returns all agents to their default state.

> **Note:** To enable or disable a single agent, use instead the commands agent on or agent off.

setup adaptiveresponse on|off|default

Command Change History

Release	Modification
11.5.1	Command introduced

setup bcaaa cache

This command is useful for troubleshooting issues with the user awareness feature. You can also use this command to see the lists of user and group names being monitored.

```
setup bcaaa cache stats show
```

Example:

```
setup bcaaa cache stats show

IP->User cache:
    Max # of nodes    : 300
    Current # of nodes: 50
    Collisions        : 0
    Positive cache hit: 0
    Negative cache hit: 2765
    Pending cache hit : 76
```

```
    Cache miss        : 56
    Insertion failure : 0
    kmalloc failure   : 0

User table:
    Max # of nodes    : 25000
    Current # of nodes: 4
    Collisions        : 0
    Cache hit         : 0
    Cache miss        : 0
    Insertion failure : 0
    kmalloc failure   : 0

Groups Of Interest table:
    Max # of nodes    : 50
    Current # of nodes: 3
    Insertion failure : 0
    kmalloc failure   : 0

BCAAA task queue:
    Queue being full  : 0
*************************************************************
* Do you want to display the contents of the following tables *
* ----------------------------------------------------------- *
* 1) IP-User Table                                           *
* 2) User Table                                              *
* 3) Groups Of Interest Table                                *
* 4) All                                                     *
* 5) Exit                                                    *
*************************************************************
```

To get more specific information about the cache, users, and groups, choose one of the options on the menu shown at the end of the output. For example, to see which group names are in the Groups of Interest table, select option 3.

Command Change History

Release	Modification
9.2.2	Command introduced

setup bcaaa force-register

Force synchronization of PacketShaper's groups of interest table with BCAAA; use this command if BCAAA was down or the PacketShaper was not connected to BCAAA while user group classes were being added or removed on PacketShaper.

```
setup bcaaa force-register
```

Command Change History

Release	Modification
9.2.2	Command introduced

setup bcaaa server-test

Verify BCAAA servers are configured properly. After you have configured the BCAAA server(s), you can verify the server configuration by testing whether an IP address returns the correct user name and group names listed in Active Directory.

```
setup bcaaa server-test <ip-address>
```

where *<ip-address>* is a single IPv4 address that you know is assigned to an Active Directory user.

Examples:

If the BCAAA servers are configured properly and the IP address is associated with a user name in Active Directory, the **test** command will return a user name for the IP address that is specified, as well as the names of the groups the user belongs to.

> **Note:** The only group names that display are ones for which a user group class has been created. The user may belong to other Active Directory groups, but they will not be listed for the user unless a class exists for that group. If the user isn't a member of any groups for which group classes have been created, NOT AVAILABLE appears for the group name.

```
# setup bcaaa server-test 10.9.116.215

 user name: BCAAA\ADMINISTRATOR
 group name(s): group-sanjose, group-engineering
```

> **Note:** If your BCAAA server was configured incorrectly (such as an invalid IP address), it may take several minutes for the *Cannot establish a connection to the BCAAA server* message to return. It may appear that the session has hung, but the prompt will appear after a few minutes.

If there is a problem with the configuration or BCAAA cannot locate the IP address in Active Directory, you will see one of the following error messages in response to the **server-test** command:

Message	Action to Take
BCAAA must be enabled in order to test the configuration.	Enable BCAAA: **setup bcaaa service on**
There is no primary server configured. Please configure a primary server and try again.	Configure BCAAA server: **setup bcaaa service primary *<host>* [*<port>*]**
Can't establish a connection to the BCAAA server. Please make sure you have configured the primary or secondary server properly.	Show BCAAA settings and make sure they are correct: **setup bcaaa show**

Message	Action to Take
The user name could not be determined for this IP: x.x.x.x	Find a valid IP address and issue **setup bcaaa test *<ip-address>*** again.

Command Change History

Release	Modification
9.2.2	User group support added
9.2.1	Command introduced

setup bcaaa service

Configure PacketShaper as a Symantec Authentication & Authorization Agent (BCAAA) client; this feature enables the PacketShaper to classify and report on user names and group names in Active Directory.

To enable or disable the BCAAA service, use:

```
setup bcaaa service on|off|default
```

To configure the settings of the BCAAA servers, use:

```
setup bcaaa service primary|secondary {<host> [<port>]}|delete
```

primary|secondary	Enter the literal **primary** or **secondary** to indicate which server you are defining. (The secondary server is used when the primary server isn t accessible.)
<host>	The IPv4 address or DNS name of the BCAAA server
[*<port>*]	The port number to access the server; if omitted, the default port 16101is used.
delete	Deletes the configuration of the primary or secondary server (whichever is specified)

Example:

This example defines a primary BCAAA server at 10.10.10.10 using the default port and a secondary server at 10.10.20.10 using port 903. The third command line enables the BCAAA service.

```
setup bcaaa service primary 10.10.10.10

setup bcaaa service secondary 10.10.20.10 903

setup bcaaa service on
```

Command Change History

Release	Modification
9.2.1	Command introduced

setup bcaaa show

Display current Symantec Authentication & Authorization Agent (BCAAA) settings and the status of each BCAAA server.

```
setup bcaaa show
```

Example output:

```
  BCAAA Setup values:
     BCAAA Service   : on
     Timeout         : 10 seconds

  BCAAA Servers:
     Type      Host                                    Port      Status
     Primary   10.9.112.240                            16101     In use
     Secondary 10.9.112.245                            16101     Available
```

The Status column has the following possible values:

Status	Description
Attempting to connect	PacketShaper is in the process of trying to establish a connection to this BCAAA server.
In use	PacketShaper is connected to this BCAAA server and is using it to look up users.
Available	PacketShaper is able to connect to this server but is using another working server (one with an "In use" status) for user lookups.
Failed connection	PacketShaper tried to connect to this BCAAA server but the connection failed. Possible reasons for a failed connection include: ■ An incorrect IP address is specified. ■ BCAAA is configured to use a different port than PacketShaper. ■ BCAAA server is down. ■ Link is down. ■ No route exists between the networks.

Issue the **setup bcaaa show** command to refresh the status.

Command Change History

Release	Modification
9.2.1	Command introduced

setup bcaaa timeout

Specify the number of seconds that PacketShaper will wait for a response from the BCAAA server when looking up a user name.

```
setup bcaaa timeout <seconds>|default
```

where *<seconds>* is a value between 2-30 seconds. The default timeout is 10 seconds.

Example:

```
setup bcaaa timeout 20
```

Command Change History

Release	Modification
9.2.1	Command introduced

setup bypass

Configure the bypass behavior when a PacketShaper is powered off or in bypass mode. When bypass is closed (the default), relays close on a powered-off PacketShaper to allow user traffic to pass unhindered between the interface ports. When bypass is open, the PacketShaper becomes a point of failure in the traffic path; this setting is required when using the standby feature. The **setup bypass** command simulates the opening and closing of relays.

```
setup bypass open|close <device>|all
 setup bypass show
```

where *<device>* is the name of the interface pair. Use the **show** parameter to list the bypass state for each device on the PacketShaper.

Traffic will only be blocked for a few seconds while the appliance is resetting or rebooting.

The *<device>* name for two-port NICs:

- slot1
- slot2 (built-in NIC on PS-S200 and PS-S400)
- slot3
- slot4

> **Note:** Slots vary by PacketShaper model.

The *<device>* name for four-port NICs:

- slot3_pair1
- slot3_pair2
- slot4_pair1
- slot4_pair2

Example:

```
setup bypass show
Device           Bypass State (power-off)
-----------------------------------------------
Slot1            closed (will engage)
Slot2            Not Bypass capable
Device           Bypass State (current)
-----------------------------------------------
Slot1            open (not engaged)
Slot2            Not Bypass capable
```

> **Note:** This command only applies to devices that use software to change the power-off bypass state. Note that some devices use physical jumpers and some do not support bypass.

Command Change History

Release	Modification
11.1.1	Updated to support PS-Series device names
8.6.3	**setup bypass** command introduced

setup capture

Capture PacketShaper's configuration. The output can be created in a portable form (without specific IP addresses) or complete form (IP addresses included). The output file includes commands to recreate the configuration.

> **Note:** This command is not intended to be a substitute for backing up PacketShaper configurations.

```
setup capture [[portable|complete] [<filename>]]
```

[portable | complete]	Indicate the format of the output. The portable file "comments out" the unit-specific details, such as the setup commands for the IP addresses. The complete file contains address information.
<filename>	The filename is saved in the 9.256/cmd folder.

Additional Information

- To restore the captured configuration, use the run command (for example, **run config.cmd**).
- Auto-discovered port classes are not recreated when you run the CMD file created with the **setup capture complete** command.
- For security reasons, the **setup capture** command does not store the look and touch passwords.

setup date

View or set the date and/or time. When initially setting the date and time, use **setup timezone**.

```
setup date [<yyyymmddhhmm>[<.ss>]]
```

Note that this command has the same functionality as the **date** command.

> **Note:** You should always do a system reset immediately after changing the date so that the underlying time-sensitive scheduled operations of the PacketShaper can be correctly initialized.

setup discover

Turn traffic discovery on or off for the default inbound and outbound classes.

When you turn on traffic discovery, the PacketShaper software monitors the traffic going through the unit and classifies the traffic by service type. The traffic discovery process inserts classes into your traffic tree.

```
setup discover off|on
```

Use **class services** to list supported protocols and services.

Use **class discover** to enable/disable traffic discovery within a specific class.

setup dns

Configure one or more DNS servers for PacketShaper to access.

```
setup dns none|default|<ipaddress> ...
```

Specify up to three IP addresses, separating each with a space, or use **none** to clear previously set addresses.

setup domain

Define a default domain name that PacketShaper can append to domain name lookups that are not fully qualified.

```
setup domain none|default|<domain_name>
```

setup email

Configure email settings for use with the scheduled command, adaptive response action files, and/or user event features.

```
setup email <address>[:<port>] [<sender name>] | none | default
```

<address>	Specify the email server using either its DNS name or IP address.
[:*<port>*]	The default SMTP port is port 25. To specify a non-standard port for email messages, enter number.
[*<sender name>*]	The *<sender name>* will appear in the From line of any email message that the user event or scheduled command feature sends out.Specify a complete mail address, including the domain name for example, *john_doe@example.com* **Note:** In the prompting mode for this command, full form can also be used for the sender that is, a quoted name followed by the explicit email address. For example: *"Bob" <bob@examplecompany.com>*

After PacketShaper has been configured to send email using the **setup email** command, use **setup show** to view the configuration.

To clear the email settings:

```
setup email none
```

If you are using PolicyCenter, use the **default** option to remove the local override. This command allows the child configuration to inherit the parent's email setting:

```
setup email default
```

Command Change History

Release	Modification
11.5.1	Command introduced

setup flowrecords engineID/engineType

Assign an identifying number to a PacketShaper, when using the NetFlow-5 flow detail record (FDR) format. EngineID and EngineType are two of the fields in NetFlow-5 headers; you can use either, or a combination, of these fields to identify the PacketShaper that is emitting records. These fields are not relevant for the Packeteer-1 and Packeteer-2 formats.

```
setup flowrecords engineID|engineType <value>|none|default
```

<value>	An integer (0-255) that identifies the PacketShaper. The default value for engineID and engineType is 0.
none	Clears the EngineID or EngineType field
default	Removes the local settings for the EngineID or EngineType field so that the unit inherits the settings of the parent configuration. If the parent configuration doesn't have any EngineID or EngineType settings, the local settings will be cleared so that the unit can inherit any future settings that are set. The **default** option is only applicable to shared mode with PolicyCenter.

Example

To assign an ID of 12 to the current PacketShaper:

```
setup flowrecords engineID 12
```

All records emitted from this PacketShaper to all defined NetFlow-5 collectors will have a value of 12 in the header's engineID field. Thus, the source PacketShaper is easily identifiable when interpreting the flow detail records.

setup flowrecords filters

Defines an include or exclude list to specify whether or not flow detail records (FDR) are emitted for traffic that matches the specified classes, services, and/or subnets. You can use this command in either of two ways:

- To define an **include** list so that flow detail records are *always* emitted for traffic that matches the specified classes, services, and/or subnets. This approach is recommended when you want flow detail records for only a few specific classes, services, and/or subnets. Flow detail records will not be emitted for traffic that does not match the class, service, or subnets specified on the include list.

  or

- To define an **exclude** list so that flow detail records are *never* emitted for traffic that matches the specified classes, services, and/or subnets. This approach is recommended when you want flow detail records for all traffic except that which matches the classes, services, and/or subnets on the exclude list.

```
setup flowrecords filters [add|remove|show] [class|service|subnet] include|exclude
[[<class name>|<class id>]|[<service name>]|[<ip:netmask>|<ip/netmask>|<ip>]]
```

[add|remove|show]	**add** adds the specified class, service, or subnet to the include or exclude list **remove** removes the specified class, service, or subnet from the include or exclude list **show** displays the classes, services, and subnets on both the include and exclude lists
[class|service|subnet]	Indicates that the FDR filter applies to a **class**, **service**, or **subnet**
include|exclude	**include** specifies that flow detail records will always be emitted for traffic that matches the class, service, and/or subnet **exclude** specifies that flow detail records will never be emitted for traffic that matches the class, service, and/or subnet
<*class name*>|<*class id*>	The name of the class or the class id
<*service name*>	The name of the service
<*ip:netmask*>	The IP address and subnet mask subject to the IP filter, where *netmask* is the subnet mask in decimal notation
<*ip/netmask*>	The IP address and subnet mask subject to the IP filter, where *netmask* is an integer (the CIDR value) that specifies the number of binary 1s in a mask
<*ip*>	The IP address subject to the IP filter

Examples

When you add classes to the include list, FDRs are emitted only for those classes. To add a class to the include list:

```
setup flowrecords filters add class include /Inbound/SNMP
```

To add the FTP service to the include list:

```
setup flowrecords filters add service include ftp
```

To add subnets specified by IP address in decimal notation to the include list:

```
setup flowrecords filters add subnet include 10.10.10.01:255.255.255.255
```

To add subnets specified by IP address and CIDR value to the include list:

```
setup flowrecords filters add subnet include 10.10.10.01/32
```

To add an IP address and all of its subnets to the include list:

```
setup flowrecords filters add subnet include 10.10.10.01
```

When you add classes, services, and subnets to the exclude list, FDRs are emitted for all classes, services, and subnets except for those specified on the exclude list. For example, to add a class to the exclude list:

```
setup flowrecords filters add class exclude /Inbound/SNMP
```

To remove a class, service, or subnet from the include or exclude list, use the remove keyword. For example:

```
setup flowrecords filters remove class include /Inbound/SNMP
```

```
setup flowrecords filters remove class exclude /Inbound/SNMP

setup flowrecords filters remove service include ftp

setup flowrecords filters remove subnet include 10.10.10.01
```

To show all FDR filters:

```
setup flowrecords filters show
```

Command Change History

Release	Modification
8.2.0	Integrated command into PacketShaper 8.x., and extended functionality so that the services and subnets can be added to FDR filter lists.
7.5.0	Command introduced

setup flowrecords id

Define the settings for a flow detail record (FDR) collector. Up to four collectors can be defined.

```
setup flowrecords id [<ID> <collectorDefinition>|off|on|none]
```

<ID>

Identifying number of the collector (1, 2, 3, or 4)

<collectorDefinition>

where *<collectorDefinition>* is

<recordType> *<ipaddr>* [*<port>* on|off]

<recordType> is the type of record format to be emitted (**netflow-5**, **packeteer-1**, or **packeteer-2**)

<ipaddr> is the IP address of the collector

<port> is the UDP port number of the collector (default = 9800)

off|on|none

on enables the collector. When a collector is enabled, PacketShaper will emit flow detail records to the collector.

off disables the collector; flow detail records will not be emitted.

none clears the collector settings; the row will be empty in the **setup flowrecords show** output.

You can enable/disable a collector when you are defining it:

```
setup flowrecords id 1 netflow-5 10.10.10.10 9800 on
```

or, after a collector has been defined:

```
setup flowrecords id 1 off
```

A collector is defined by its record type (NetFlow-5, Packeteer-1, or Packeteer-2) and its location (IP address and UDP port number). You can define collectors with the same IP address but different record types, or with the same record type but different IP address. For example, you can create two collectors with the same IP address (but different ports), with one collector collecting NetFlow data records and the other collecting Packeteer-2 data records.

To view your collector settings, use the setup flowrecords show command.

Examples

To define a collector that collects Packeteer-2 flow detail records:

```
setup flowrecords id 1 packeteer-2 10.10.10.1 9800 on
```

Because 9800 is the default port and "on" is the default, you can use the following alternative command:

```
setup flowrecords id 1 packeteer-2 10.10.10.1
```

To turn off collector 1 (assuming collector 1 has been previously defined):

```
setup flowrecords id 1 off
```

With the above command, PacketShaper will stop emitting flow detail records to collector 1, but will retain the collector settings. To start emitting records again, use this command:

```
setup flowrecords id 1 on
```

To clear the settings for collector 3:

```
setup flowrecords id 3 none
```

setup flowrecords show

Display the flow detail record (FDR) collector settings. Use this command to see the collectors that have been configured and check which ones have been enabled. This command is also useful to look up the ID number associated with a collector (the ID is needed for defining and clearing collector settings).

```
setup flowrecords show

ID  RecordType      CollectorIP    Port   Enabled
 1   packeteer-2     10.10.10.1     9800   on
 2   netflow-5       10.10.10.2     9800   off
 3
 4
```

In the above sample output, two collectors have been defined. The first collector (ID of 1) collects Packeteer-2 flow detail records and is currently enabled. The second collector (ID of 2) collects NetFlow-5 records but is currently disabled. Collector IDs 3 and 4 have not been defined.

setup gateway

Configure a gateway to handle network operations initiated *from* the unit. For example, a ping requires a gateway for non-local routing.

```
setup gateway <ipaddress> | none
```

Specify **none** if there isn't a gateway or to clear the gateway setting.

setup guide

Execute the automatic setup feature to configure the unit.

```
setup guide
```

For Guided Setup details, see PacketGuide .

setup heartbeat

Configure the PacketShaper to emit messages (heartbeats) to the Symantec heartbeat server. Using the information contained in the heartbeat messages, Symantec is able to compile statistics on the stability of various software releases and hardware products. The heartbeats can also be used to identify and resolve defects.

```
setup heartbeat on|off|default|show
```

Heartbeat emission is enabled by default. Symantec recommends that you *not* disable the feature. Be assured that the messages are encrypted and sent securely via HTTPS. The size of the daily heartbeat message is negligible (30-40K) and has virtually no impact on PacketShaper performance.

Examples:

To disable the heartbeat feature:

```
setup heartbeat off
```

To find out when/if the last heartbeat message was sent:

```
setup heartbeat show
```

```
    Heartbeat                  : On
    Most Recent Attempt Status : Heartbeat sent successfully, Jun 17 2010 10:31:06
```

If PacketShaper isn't able to send the heartbeat, the **setup heartbeat show** output will indicate the reason for the failure:

- DNS resolution errors
- Transport protocol error
- File I/O errors
- Error running certain commands
- Insufficient free hard disk space

> **Note:** Depending on how long it takes for the configuration system to initialize DNS server entries, the **setup heartbeat show** output may show a DNS resolution error if the PacketShaper cannot resolve the heartbeat server name via DNS quickly enough.

Command Change History

Release	Modification
8.5.1	**setup heartbeat** command introduced

setup images active

Select the image you want to run. After rebooting, this image will be the one running on the system.

```
setup images active <release_id>
```

Use the setup images show command to see the release ID number to specify. Use the version command to see the release ID of the current active image.

You will need to reset the PacketShaper to complete the activation process.

> **Note:** If connected to the serial console at boot up time, you can press the space bar and select the desired image to boot.

Command Change History

Release	Modification
11.1.1	Command introduced

setup images add

Add an image to the PacketShaper. Use this command so that you can upgrade to a new release. The PacketShaper can manage up to six images. If your PacketShaper already has six images installed and you add another image, the oldest unlocked image will be replaced with the new image.

```
setup images add <URL to image> | file://<filename>
```

where:

<URL to image> is the path to an image on a web server that the PacketShaper has access to. Before initiating the **setup images add** command, you need to log in to BTO, and download the new software image to a web server the PacketShaper can access.

file://*<filename>* is the filename of an image on the PacketShaper. Before initiating the **setup images add** command, you need to log in to BTO, download the new software image to a local workstation or server, and then upload it to the 9.258/image folder on the PacketShaper.

> **Note:** Make note of the new filename as you will need to specify this name when adding the image.

Examples:

```
setup images add http://webserver.mycompany.com/images/542386.bcs
```

```
setup images add file://542386.bcs
```

Additional Information

- When you add an image, the image manager assigns it a unique release ID (such as 532537); you must refer to this ID when activating, removing, locking, and unlocking images.
- The **setup images add** command does not support https.

Command Change History

Release	Modification
11.1.1	Command introduced

setup images detail

Display details about an image, such as the software version, whether the image is locked, and when the image was last booted.

```
setup images detail <release_id>
```

Use the setup images show command to see the release ID numbers.

Example:

```
#setup images detail 133755
--------------------------------------------------------------------------------
Get Image Detail: 133755
--------------------------------------------------------------------------------
  Release Id       : 133755
  DisplayName      : PacketShaper 11.1.1.1, Release ID: 133755
  Product Name     : PacketShaper
  Version          : 11.1.1.1
  Default Image    : YES
  Image Locked     : NO
  Last Boot Failed : NO
  Image Booted     : YES
  Boot Timestamp   : Mon Feb  3 14:35:04 2014
```

Command Change History

Release	Modification
11.1.1	Command introduced

setup images lock

Lock an image to protect it from accidental deletion.

```
setup images lock <release_id>
```

Use the setup images show command to see the release ID numbers.

If you later decide you want to delete a locked image, you must use the setup images unlock command first.

Command Change History

Release	Modification
11.1.1	Command introduced

setup images remove

Remove an image from the PacketShaper.

```
setup images remove <release_id>
```

Use the setup images show command to see the release ID numbers.

You cannot remove a locked image or the active image. If an image is locked, use the setup images unlock command before issuing the **setup images remove** command.

Command Change History

Release	Modification
11.1.1	Command introduced

setup images show

Display a list of installed images and their version and release ID numbers. You will need to refer to the release ID when activating, removing, locking, and unlocking images.

```
setup images show
```

Example:

```
setup images show
--------------------------------------------------------------------------------
List of Installed Images
--------------------------------------------------------------------------------
PacketShaper 11.2.1.1  Release Id: 533930 (A)
PacketShaper 11.1.15.1 Release Id: 534295
PacketShaper 11.1.14.1 Release Id: 534667
PacketShaper 11.1.13.1 Release Id: 535202
Total 4 images  (A) Active image
--------------------------------------------------------------------------------
```

The (A) marker indicates which image will be loaded when the PacketShaper is next rebooted. This may or may not be the currently running image; to see which image is currently active, use the version command.

Command Change History

Release	Modification
11.2.1	PacketShaper OS version number added to output
11.1.1	Command introduced

setup images unlock

Unlock an image that you no longer want to protect from deletion. You have to unlock a locked image before you can remove it.

```
setup images unlock <release_id>
```

Use the setup images show command to see the release ID numbers.

Command Change History

Release	Modification
11.1.1	Command introduced

setup ipaddress

Update the unit's IP address and subnet mask. Use dotted-decimal address notation for both the IP address and net mask for example, *10.10.10.10*.

```
setup ipaddress <addr> <netmask>
```

You will need to reset the PacketShaper after changing its management IP address. PacketShaper will not send FDRs and link settings may not be accurate until you reset the unit.

setup keys show

Display license keys along with their activation and expiration dates.

```
setup keys show
```

Command Change History

Release	Modification
11.1.1	Command introduced

setup link

Configure the access link capacity. To effectively manage the traffic on the link, PacketShaper must know the capacity it is managing.

> **Note:** PacketShaper will enforce the link size that you set.

```
setup link inbound|outbound|default [<size_bps>|default]
```

Specify a rate as either a bits-per-second value or a symbolic name, as shown in the following list of valid link sizes.

<n>	Size in bits per second
*<n>***k**	Size in kilobits per second
*<n>***m**	Size in megabits per second
*<n>***g**	Size in gigabits per second
T1	1.5 Mbps
E1	2 Mbps
T3	45 Mbps

Examples:

```
setup link inbound 1500000

 setup link outbound 1.5m

 setup link inbound T1
```

Considerations

- For full-duplex Ethernet, enter the total link speed for the inbound and outbound rates. Because full-duplex has wires that can simultaneously communicate in both inbound and outbound directions, you should enter the same rate for Inbound Rate and Outbound Rate. For example, if you have two T1 lines (3 Mbps), you should enter 3M for Inbound Rate and 3M for Outbound Rate.

- For half-duplex Ethernet, split the rate between the inbound and outbound links. For example, if you are managing 10 Mbps Ethernet, you could configure 5 Mbps for the inbound rate and 5 Mbps for the outbound rate.

- If your appliance is using NICs to manage different WAN links and you don't want to control each NIC separately, the rate should be the size of the smallest NIC. For example, if you have two 100 Mbps NICs managing two links, you should specify 100M for the rate.

On the other hand, if you want to control each link separately, the rate should be the sum of the link speeds on all devices. For example, if the built-in device is controlling a T1 line (1.5 Mbps) and a NIC option card is managing two T1 lines (3.0 Mbps), you should specify 4.5M for the rate. To control traffic across each link separately, you can create a class for each device (for example, *Slot1* and *Slot3*) and assign partitions that match the link size (1.5M for the Slot1 class and 3.0M for the Slot3 class).

If your appliance is using two NICs to manage a single WAN link, specify the WAN link speed for the rate. Although the Info page will give you an error message (such as Link speed of 155 Mbps exceeds outside NIC speed of 100 Mbps) in the latter situation, it is still appropriate to specify the actual size of the link for the rate.

- Software configuration determines maximum shaping capacity.

> **Note:** 10BaseT links rarely reach the 10 Mbps limit. Keep Ethernet's practical limits in mind when configuring rates.

setup load

Install a new configuration file and reboot to activate the configuration. This installation replaces the cfg/basic.cfg file.

```
setup load <path>
```

Specify the explicit file pathname.

> **Note:** To load the traffic configuration and sharable configuration settings (such as passwords, SNTP, and Syslog), use the class load command.

setup loadshedding

Configure the load shedding feature. This feature prevents the PacketShaper from being overloaded with packets due to viruses or attacks that spew out a high volume of traffic (such as ICMP or DCOM). Note that this feature is not designed to block all traffic from infected or misbehaving clients; it is designed to shed just enough traffic to keep the PacketShaper out of an overload condition. It allows enough of the "bad" traffic through so that you can use PacketShaper diagnostic tools (traffic history, hostdb info, policy flowlimit, packetcapture) to analyze and contain the problem.

```
setup loadshedding enable|disable
 new on|off
 clientFPM|serverFPM|failedFPM|TCBConn|UCBConn <value>|default
 exception add|del list:<hostlist>|<ip-addr>|<dnsname>|all
 show
```

where

enable|disable	Enable or disable the load shedding feature. Default is **disable**.
new on|off	Drop packets of new flows only (**on**) or drop packets regardless of whether they are new or existing (**off**). Default is **on** (but is not in effect unless load shedding is enabled).
clientFPM|serverFPM|failedFPM|TCBConn|UCBConn *<value>*|default	Adjust the parameters for new flows per minute (client, server, failed) or connections (TCP or UDP). If the parameters are not specified, the default values are used. *New flows per minute* is the rate of initiation of new flows from a host (client) or to a host (server). *TCBConn* is the number of active TCP flows that a particular host has at a given time. *UCBConn* is the number of non-TCP flows a particular host has at a given time. *Failed flows* are TCP flows that do not establish a complete connection, such as TCP connection requests from a SYN flood attack. The *<value>* parameter can be any whole number from 10 to 1,000,000 (inclusive) for *clientFPM*, *serverFPM* and *failedFPM*. The <value> parameter for *TCBConn* and *UCBConn* can be any whole number from 0 to 1,000,000, inclusive.
exception add|del list: *<hostlist>*| *<ip-addr>*|*<dnsname>*|all	Add the host list name, IP address, or DNS name of a host to be excluded from load shedding (packets will not be dropped from these servers). To specify multiple hosts, use a space between each one. Do not use subnets or ranges of addresses to specify hosts. For example: `setup loadshedding exception add 10.1.1.1 172.19.5.6 olympia` Use **exception del** to remove a host that you have previously added to the host exception list, or **exception del all** to remove all hosts. (**Note:** The **all** parameter is used with **del** only.)
show	Display load shedding settings.

Load shedding is disabled by default. To enable load shedding:

```
setup loadshedding enable
```

The load shedding feature drops packets intelligently when the PacketShaper sees an excessive amount of traffic. For client flows, packets will be dropped as the unit approaches its load capacity and when both of the following conditions are true:

- New flows per minute for the client exceeds the *clientFPM* value, AND
- The number of TCP flows for the client exceeds the *TCBConn* value OR the number of UDP flows for the client exceeds the *UCBConn* value

Load shedding works similarly for server flows per minute. Packets will be dropped as the unit approaches its load capacity and when both of these conditions are true:

- New flows per minute for the server exceeds the *serverFPM* value, AND
- The number of TCP flows for the server exceeds the *TCBConn* value OR the number of UDP flows for the server exceeds the *UCBConn* value

For failed flows, packets will be dropped as the unit approaches its load capacity and failed flows per minute for the client or server exceeds the failedFPM value.

> **Note:** To see the current values of each host's new flows per minute, use the hostdb info command.

By default, load shedding will drop packets of new flows only existing flows will not be dropped. If you want to remove this limitation, use the following command:

```
setup loadshedding new off
```

Use the **setup loadshedding show** command to display the current, default, minimum, and maximum parameters for load shedding.

```
Load Shedding:                  Disabled
Shed New Flows Only:            Enabled
=====================================================================
Load Shedding Parameters        Current   Default   Min       Max
=====================================================================
Client FPM                      1000000   1000000   10        1000000
Server FPM                      1000000   1000000   10        1000000
Failed FPM                      1000000   1000000   10        1000000
TCBConn                         100       100       0         1000000
UCBConn                         100       100       0         1000000

Load Shedding Host Exception List:
179.21.1.3      server2.test.com
10.1.1.1        main.test.com
10.1.1.2        server1.test.com
```

If you have certain hosts that you want to exclude from load shedding (for example, you don't want load shedding to drop packets from DNS servers), you can create a host exception list. You can either:

- Create a host list with the hl new command and then specify the host list with the **setup loadshedding exception add list:*<hostlist>*** command.

  or

- Add the hosts individually with the **setup loadshedding exception add *<ip-addr>*|*<dnsname>*** command.

To see which hosts have exceeded the load shedding thresholds and have flows being shed, use the hostdb info command. A "+" next to the New Flows Per Minute value for Client, Server, or Failed indicates load shedding is occurring or has recently occurred.

Command Change History

Release	Modification
8.3.1	The minimum value for *TCBConn* (active TCP flows) and *UCBConn* (active non-TCP flows) changed to 0 flows. Previous versions of PacketShaper required a minimum value of 5 flows.
7.0.0	command introduced

setup message

Configure a message that will display before logging into the PacketShaper. The message displays before you log in via the browser login page, before logging in using a remote login utility (such as putty), and when you first console connect to the unit. This feature is useful for informing users about the company's access policies and consequences for unauthorized use.

```
setup message {set <message>}|show|default
```

where

set *<message>*	Defines the message text. The text should be enclosed in quotation marks and can be up to 511 characters long.
show	Displays the content of the login message
default	Clears the message text. In PolicyCenter's shared mode, the unit will then be able to inherit the message of the parent configuration.

Examples

```
setup message set "Access to this system is restricted to authorized users only."
Message set to: "Access to this system is restricted to authorized users onl...

setup message show

Configured Message:
Access to this system is restricted to authorized users only.
```

Notes

- Quotation marks indicate the beginning and end of the login message. You cannot use a quotation mark within the body of the login message.
- If you want to display a message that is longer than 511 characters, you can create a text file that contains your message text. Name the file login.txt and upload it to the 9.256/ folder. The first 2048 characters of the text file will display after any message that is configured with the **setup message set** command. Thus, the text file is appended to the message text, allowing the message to have a total approximate length of 2500 characters. Note that quotation marks are allowed in the login.txt file.
- The **setup message show** command does not display the content of the login.txt file.
- The message can be configured in the browser interface as well.

setup nic

Set the PacketShaper's speed and duplex state.

```
setup nic <device> auto|(10bt|100bt) (full|half)|(1000b|10000b) full
```

where *<device>* is the interface name:

Dual-port NICs:

- slot1_in
- slot1_out
- slot2_in
- slot2_out
- slot3_in
- slot3_out
- slot4_in
- slot4_out

Quad-port NICs:

- slot3_in1
- slot3_out1
- slot3_in2
- slot3_out2
- slot4_in1
- slot4_out1
- slot4_in2
- slot4_out2

> **Note:** Slots vary by PacketShaper model.

Management interface:

- mgmt

Specify **auto** (auto-negotiate) to automatically configure the unit for the appropriate mode. If you do not specify a state, it defaults to **auto**.

Additional Information

- Whenever you wish to change Network Interface Card (NIC) settings, always select auto-negotiate first, then select a different value if desired. Do not change from one non-auto setting to another non-auto setting directly; re-negotiation may fail.

- Although you can specify different fixed speeds on the Inside and Outside interfaces, such a configuration will result in a network interruption if the PacketShaper is turned off because the end devices will not be able to negotiate the correct speed for the link.

- PacketShaper does not support a 1000b or 10000b half-duplex interface.

Command Change History

Release	Modification
11.1.1	Updated to support PS-Series device names

setup password

Configure a touch (read/write) or look (read-only) password. The default password for the touch user is *touch* and the default password for the look user is *look*.

```
setup password look|touch
```

Passwords can be up to nineteen characters long and are case sensitive. They can consist of a combination of letters, numbers, and all special characters on the U.S. keyboard. Passwords cannot be blank. Passwords are set to their defaults when the **setup reset all** or **setup reset clear** commands are executed.

You will be prompted to enter the old password, type a new password, and retype the password to confirm. For example:

```
setup password touch

Old touch password: (none)
New touch password:
Confirm touch password:
Changed the touch password
```

To abort this command and return to the command prompt, press Ctrl+D.

Command Change History

Release	Modification
11.1.1	New default passwords for look and touch users

setup portal ip

Assign a second IP address to the customer portal. This allows customers to directly display the customer portal login page using this address instead of the http://x.x.x.x/customer URL.

```
setup portal ip <address> [<mask>]
```

where *<address>* is the IP address assigned to the customer portal and *<mask>* is the subnet mask for the network where the unit resides. The address must be on a different subnet from the main portal address.

To clear the address, use:

```
setup portal ip none
```

Command Change History

Release	Modification
11.6.1	Command introduced

setup portal message

Set the system-wide message-of-the-day for the customer portal feature. This message is displayed if the customer doesn t have a specific message configured with the **portal new** command.

```
setup portal message <new message>
```

The message can be up to 128 characters long and must be enclosed in quotes if it contains spaces.

Command Change History

Release	Modification
11.6.1	Command introduced

setup portal show

Lists the system-wide message-of-the-day.

```
setup portal show
```

Command Change History

Release	Modification
11.6.1	Command introduced

setup radius acct

Set up or change the configuration of the RADIUS accounting service. This feature allows you to have an audit trail for user logins.

To define the RADIUS accounting service, use:

```
setup radius acct primary|secondary <host> <shared_secret> [<port>] |delete|override
```

primary|secondary	Enter the literal primary or secondary to indicate which server you are defining. (The secondary server is used when the primary server isn t accessible.)
<host>	The IP address or DNS name of the RADIUS accounting server
<shared_secret>	The designated secret for the server; quotes are not required
[*<port>*]	The port number to access the server; if omitted, the default port is used
delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	Inherits setup of the primary or secondary server from PolicyCenter

To turn the service on or off, or to return the service to its default **off** value, use:

```
setup radius acct on|off|default
```

Example:

```
setup radius acct primary 10.10.10.10 bobolink

setup radius acct secondary 10.10.20.10 parrot

setup radius acct on
```

This example defines a primary accounting server at 10.10.10.10 which has a shared secret of bobolink, as well as a secondary server at 10.10.20.10. The third command line enables RADIUS accounting service. Once this service is configured and enabled, PacketShaper will send a PW_STATUS_START accounting message to the accounting server when a user logs in and a PW_STATUS_STOP message when a user logs off or is disconnected.

Command Change History

Release	Modification
11.5.1	Command introduced

setup radius auth

Set up or change the configuration of the RADIUS authentication service. RADIUS authentication is an optional method for users to log into the PacketShaper browser interface, command-line interface, or customer portal or when FTPing to the unit. Using third-party RADIUS servers enables you to have central configuration of user accounts.

```
setup radius auth primary|secondary <host> <shared_secret> [<port>] |delete|override
```

primary|secondary	Enter the literal primary or secondary to indicate which server you are defining. (Note: The RADIUS client uses the secondary server when the primary server isn t accessible or authentication failed.)
<host>	The IP address or DNS name of the RADIUS authentication server
<shared_secret>	The designated secret for the server; quotes are not required
[*<port>*]	The port number to access the server; if omitted, the default port is used

delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	Inherits setup of the primary or secondary server from PolicyCenter

To turn the service on or off, or to return the service to its default **off** value, use:

```
setup radius auth on|off|default
```

Example:

```
setup radius auth primary 10.10.10.10 bobolink

setup radius auth on
```

This example first defines a primary authentication server at 10.10.10.10 which has a shared secret of *bobolink*. The second command line enables RADIUS authentication service. Once this is configured and enabled, PacketShaper will prompt users for user name and password when they log into PacketShaper.

Command Change History

Release	Modification
11.5.1	Command introduced

setup radius interval

Set the amount of time for RADIUS to wait for a response from a server. By default, the RADIUS client waits 5 seconds before retrying a login when the RADIUS server fails to respond.

```
setup radius interval <seconds>|default
```

where *<seconds>* is a value between 1 and 30 seconds. For example:

```
setup radius interval 20
```

In this example, the retry interval is 20 seconds; this interval applies to any configured RADIUS server.

To return to the default retry interval, use:

```
setup radius interval default
```

Command Change History

Release	Modification
11.5.1	Command introduced

setup radius limit

Set the number of retry attempts the RADIUS client will make to a server before cancelling the login. By default, if the RADIUS server fails to respond, the RADIUS client will try to log onto the server three times before reporting a server failure. If you have specified a secondary server, the RADIUS client will alternate attempts to log onto each server.

```
setup radius limit <n>|default
```

where *<n>* is a value between 1 and 10. For example:

```
setup radius limit 6
```

In this example, the RADIUS client will try to log onto the server six times.

To return to the default retry limit, use:

```
setup radius limit default
```

Command Change History

Release	Modification
11.5.1	Command introduced

setup radius method

Select the RADIUS authentication method: PAP, CHAP, or MSCHAP.

```
setup radius method PAP|CHAP|MSCHAP|default
```

PAP *(Password Authentication Protocol)*	With PAP, the user name and password are transmitted in clear, unencrypted text. ASCII or PAP authentication is required for RADIUS configurations that require access to clear text passwords (for example, when passwords are stored and maintained in a database external to the RADIUS server).
CHAP *(Challenge Handshake Authentication Protocol)*	In some environments, CHAP may be preferred for greater security. The RADIUS server sends a challenge that consists of a session ID and an arbitrary challenge string, and the user name and password are encrypted before they are sent back to the server. CHAP is the default authentication method.

MSCHAP (Microsoft Challenge Handshake Authentication Protocol)	This protocol is similar to CHAP, but with MS-CHAP authentication, the RADIUS server can store an encrypted version of a user password to validate the challenge response. Standard CHAP authentication requires that the server stores unencrypted passwords. **Note:** MS-CHAP v1 and v2 are supported. PacketShaper attempts authentication with MS-CHAP v2 first. If the remote server doesn't support v2 or if authentication is denied, PacketShaper re-attempts authentication with MS-CHAP v1. **Limitation:** PacketShaper SSH does not support the MS-CHAP or CHAP authentication methods. Regardless of which method the PacketShaper is configured to use for RADIUS authentication, PAP will be used when logging in with SSH to the CLI. Note that this limitation does not apply when logging in to the web UI: the configured method will be used during authentication.

Command Change History

Release	Modification
11.5.1	Command introduced

setup radius show

Display the current RADIUS settings. Use this command to verify that RADIUS authentication and accounting are enabled, to see the current settings for the retry limit and retry interval, and to view configuration settings on each of the RADIUS servers.

```
setup radius show
```

```
Authentication :on
 Accounting :on
 Retry limit :3
 Retry interval :5

 Service records:

 Type    Host          Port    Secret
 auth1   10.7.55.1     1812    testing123
 acct1   10.7.55.1     1813    testing123
```

Command Change History

Release	Modification
11.5.1	Command introduced

setup reset

Return to the factory-default configuration and reboot the unit.

```
setup reset [all|clear]
```

setup reset	Resets the PacketShaper settings (for example, NIC speed and IP address, but not the traffic tree) to the factory default state, and then reboots the unit
	Resets all settings and the configuration to the factory default state
setup reset all	Use **`setup reset all`** only when you want to reset *all* configuration settings basic configuration, the traffic tree with its classes, policies, and partitions; measurement data to the factory-default settings. This command deletes all custom service groups and returns any moved services back to their original groups. Passwords are reset to their defaults (*look* for the look user, *touch* for the touch user).
setup reset clear	Clears all files on the "/public" folder, all logs, RADIUS/TACACS+ user and server information, in addition to the configuration settings that **`setup reset all`** clears.

These commands remove the unit's IP address, making it unreachable on your network until it is reconfigured using Guided Setup.

To reboot the unit without modifying the settings, use the reset command.

Command Change History

Release	Modification
8.3.1	The **clear** parameter is introduced.

setup secure

Limit management access from the inside, outside, or management interface.

```
setup secure inside|outside|mgmt on|off|default|list <addr>[:<mask>]...
```

Use the **setup secure outside on** command to secure the outside interface, that is, the Internet. For example, when the outside interface is set to secure, Telnet, HTTP, FTP and ping requests from external sources will not be permitted. By default, none of the interfaces are secured.

The **list** parameter enables access to up to 16 IP addresses, separated by spaces. This is an exception list–the interface is secured except for the IP addresses on the list. To specify a subnet, use the format: *ipaddress:subnet_mask* or *ipaddress/CIDR*. The **list** option accepts IP addresses only, not host names. To find the IP address associated with a host name, use the dns lookup CLI command.

Notes:

- If you secure the interfaces, you will be able to access the unit *only* via a console connection. The browser interface will be disabled because you will not have management access over the network. Another way to secure the interface is to specify a list of IP addresses that can access the unit. For example, **setup secure outside list 10.1.1.100 10.1.12.1** would allow access from only two IP addresses.

- Keep in mind that securing an interface means that queries such as DNS and SNTP cannot be made via the secured interface. Consider using the **list** option and including these servers and your gateway in the list.

- The WebPulse classification features require access to a number of outside web servers. Therefore, do not completely secure the outside interface. Instead, use the **list** option and add the IP addresses of the WebPulse service points (use the **setup webpulse show service** command to find the IP addresses of the one or two fastest servers). In addition, you will need to specify the URLs of other servers PacketShaper needs to access. If you are using a web proxy, you will also need to add this server's IP address to the list if it will be accessible via the outside interface.

- The PacketShaper will not be able to process local ARP requests via a secured interface.

- If you secure the outside interface and your gateway is on the outside, a "gateway not found" message will be displayed in the login banner or on the info page. In this state, tasks such as upgrading the software image from a non-local address will be disabled.

setup shaping

When shaping is turned on, traffic is classified and measured, and control policies are enforced. When shaping is off, traffic is classified and measured but not managed.

```
setup shaping on|off|bypass|passthru|watch
```

Where:

on	Turns traffic shaping on
off	Turns off shaping mode (traffic, bypass, passthru, watch)
bypass	Sets the unit to pure bypass mode. Bypass mode prevents both packet shaping and further network management access; it is as if the unit were removed, and cables connected around it.
passthru	Turns off all shaping, classification, and measurement
watch	Sets the unit into a non-inline, monitor-only mode.

setup show

Display the basic configuration.

```
setup show
```

The output is divided into non-sharable (local) and sharable settings. The sharable settings are part of the configuration file (config.ldi). If a configuration is loaded on another unit, the sharable settings will be copied to the other unit (see config save and config load).

```
Non-sharable (local) settings:
  IP address:             10.9.28.170 Subnet mask: 255.255.252.0
  Gateway:                10.9.28.1
  DNS server(s):          10.2.2.100 10.2.2.11 10.2.2.10
  Default domain:         (none)
  Date, time, timezone:   Thu Nov  7 10:12:49 2013 PST (LosAngeles)
  Category Map:           v9 downloaded on Wed Nov  6 10:28:38 2013
  Web Applications Map:   downloaded on Wed Nov  6 10:28:38 2013
  Slot2_inside_0 nic speed: 1Gbps (Twisted pair) Full-Duplex(auto-negotiate)
  Slot2_outside_1 nic speed: 1Gbps (Twisted pair) Full-Duplex(auto-negotiate)
  Slot1_inside_0 nic speed: No link (manual)
  Slot1_outside_1 nic speed: No link (manual)
  Mgmt_inside_0 nic speed: 1Gbps Full-Duplex(auto-negotiate)
  Installed Keys:         demo                on (Expires: DEC 31, 2013)

Sharable settings:
  Inside interfaces:      unsecure
  Outside interfaces:     unsecure
  Look password:          ******
  Touch password:         ******
  Link speed:             1G (1000BT))
  Packet shaping:         on
  Traffic discovery:      on
  WebPulse:               on
  Web Proxy:              off
  SNTP Client:            off
  SNTP Primary Server:    time.nist.gov
  SNTP Secondary Server:  time-a.nist.gov
  SNTP Poll Seconds:      300
  BCAAA:                  on
  BCAAA Primary Server:   10.9.30.140:16101
  BCAAA Secondary Server: 10.9.30.140:16101
  BCAAA Timeout Seconds:  10
  Syslog:                 off
```

setup snmp accessgroup

Each SNMP access group is defined by a group name, a security model (and level), and a set of views that specifies which types of MIB data that access group can read or write.

There are two different commands to modify access group settings. Modify the group settings on a PacketShaper in local mode or a top-level PolicyCenter configuration with the command **setup snmp accessgroup modify**. Use the command **setup snmp accessgroup override** on a PacketShaper in shared mode or a PolicyCenter child configuration to create a local copy of a SNMP group that overrides the inherited SNMP group.

```
setup snmp accessgroup new|modify|override <groupname> [noAuthNoPriv|authNoPriv|authPriv]
[read <viewname>] [write <viewname>]}
```

where

<groupname>	The name of the user group you are creating or modifying. An access group name can be up to 32 characters; hyphens, underscores, and periods are acceptable. If the group name contains spaces, it must be enclosed in quotations marks, for example "admin group."
noAuthNoPriv authNoPriv authPriv	SNMPv3 access groups support the following security levels: ■ **noAuthNoPriv**: Identifies a user for access control, but does not provide authentication. ■ **authNoPriv**: Identifies a user for access control, and authenticates the user's password. ■ **authPriv**: Identifies a user for access control, authenticates the user's password, and provides encryption. If you do not specify a usm security model, the group will use the default **noAuthnoPriv**.
read *<viewname>*	Access groups have *read* (look) access to the information specified by the read view. To give the group read access to all MIB data, specify the predefined view name **isoAll** for the *<viewname>* parameter. To block all read access, specify **isoNone**. To limit a group's read access to a subset of available MIB data, enter the name of a user-defined view created with the **setup snmp view** command. If you do not specify a read view, the group will apply the default **isoAll** setting.
write *<viewname>*	Access groups have *write* (touch) access to the information specified by by the write view. To give the group write access to all MIB data, specify the predefined view name **isoAll** for the *<viewname>* parameter. To block all write access, specify **isoNone**. To limit a group's write access to a subset of available MIB data, enter the name of a user-defined view created with the **setup snmp view** command. If you do not specify a write view, the group will apply the default **isoAll** setting.

Examples:

```
setup snmp accessgroup new engineering usm authpriv read isoall write isoall

setup snmp accessgroup new admin usm authpriv read snmpTraps write isoNone
```

> **Tip:** Since each individual SNMP setting change requires a full configuration reload, Symantec recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

Delete an Access Group

To delete an access group, use:

```
setup snmp accessgroup delete <groupname>
```

where *<groupname>* name of the group you want to delete. Note that you will not be able to delete a group that currently has users assigned to it.

Example:

```
setup snmp accessgroup delete marketing
```

View Access Group Settings

To view current SNMP access group settings, issue the command:

```
setup snmp accessgroup show
```

Example output:

```
AccessGroupName      Model Level        ReadViewName         WriteViewName Refs Status
admin                usm   authPriv     all_mib              isoNone         10   ok
engineering          usm   authPriv     all_mib              Trap             23  ok
v2                   v2c                isoAll               isoAll           1   ok
v1_users             v1                 isoAll               isoAll           0   ok
```

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series
8.3.1	Command introduced

setup snmp destinations

The IP address(es) of SNMPv1 trap destination(s). For *<ipaddress>*, you can specify up to eight IP addresses in dotted-decimal notation, separated by spaces. Note that DNS names are not supported. If the **default** parameter is specified, all IP address destinations are cleared.

```
setup snmp destinations <ipaddress>...|default|none
```

Example:

```
setup snmp destinations 172.22.20.156 172.23.21.19
```

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series

setup snmp disable

Disable SNMP functionality on the PacketShaper. When SNMP is disabled, the PacketShaper is not able to accept incoming SNMP requests or send SNMP traps. SNMP is disabled by default.

```
setup snmp disable
```

> **Tip:** Since each individual SNMP setting change requires a full configuration reload, Symantec recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

Command Change History

Release	Modification
11.2.1	Command introduced

setup snmp enable

Enable SNMP functionality on the PacketShaper. (SNMP is disabled by default.) When SNMP is enabled, the PacketShaper is able to accept incoming SNMP requests and send SNMP traps.

```
setup snmp enable
```

For full SNMPv1 or v2 functionality on PacketShaper, you also need to configure the look or touch community strings and set trap destinations . For SNMPv3, you need to configure notify targets.

Command Change History

Release	Modification
11.2.1	Command introduced

setup snmp look|touch

Set SNMP look or touch community strings (passwords). SNMP is not functional on PacketShaper until you configure the look community string and enable it with the "setup snmp enable" on the previous page command.

```
setup snmp look|touch <string>|default|none
```

Where

Parameter	Description	With default setting
look	The context-sensitive SNMP look (read) password. Allowed characters are: [A-Z] [a-z] [0-9] [_-.,@!$%^*~=+]	Look password is cleared (set to **public**)
touch	The context-sensitive SNMP touch (read/write) password. Allowed characters are: [A-Z] [a-z] [0-9] [_-.,@!$%^*~=+]	Touch password is cleared (set to **public**)

The community string can contain alphanumeric characters, hyphens, underscores, and periods; all other special characters should be avoided.

To see the current settings for the look or touch community string values, use the setup show command.

Examples:

```
setup snmp look lookpwd1
setup snmp touch touchpwd2

setup snmp show

 SNMP config mode: simple
 SNMP look community: lookpwd1
 SNMP touch community: touchpwd2
 SNMP Trap destinations: (none)

ViewName           SubtreeOID        Type              Refs    Status
  B isoAll           iso             included    0           ok
  B isoNone          iso             excluded    0           ok
```

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series

setup snmp remoteuser

An SNMP remote user defines a user or a management system that receives notification of SNMPv3 traps and informs. Unlike a local SNMP user, a remote user is not associated with an access group and therefore has only a notify view, rather than a read or write view.

There are two different commands to modify remote user settings. Modify remote user settings on a PacketShaper in local mode or a top-level PolicyCenter configuration with the command **setup snmp remoteuser modify**. Use the command **setup**

snmp remoteuser override on a PacketShaper in shared mode or a PolicyCenter child configuration to create a local copy of a SNMP remote user that overrides the inherited SNMP remote user.

```
setup snmp remoteuser new|modify|override <username> [<engine-id>] [auth {md5|sha} <auth-pw>] [priv {des|3des|aes128|aes192|aes256} <priv-pw>]
```

where

<username>	Name of the user you are creating or modifying. Remote user names can have up to 32 characters; hyphens, underscores, and periods are acceptable. If the name includes a space, it must be enclosed within quotation marks, for example "John Doe." Each SNMP remote user name must be unique.
<EngineID>	An SNMP Engine ID identifies an SNMP engine that will receive trap and inform notifications. The default Engine ID for a remote SNMP user is *LocalSnmpId*, the SNMP agent's own SNMP Engine. If you omit this parameter, the remote user will user this default *LocalSnmpId* Engine ID. To specify a *different* remote SNMP engine with which this user can communicate, specify the 24-digit hexadecimal Engine ID of a remote SNMP engine.
auth {md5|sha}	If the remote user requires authentication, specify either the **MD5** or **SHA** authorization protocol and enter an authentication password for the user. If the remote user does not require authentication, this parameter can be omitted.
<auth-pw>	Authorization password for the user. Passwords can have up to 32 characters; hyphens, underscores, and periods are acceptable.
priv {des|3des|aes128|aes192|aes256}	Specify one of the following privacy protection protocols if the remote user requires privacy protection. Otherwise, this parameter is not required. ■ **des**: CBC-DES Symmetric Encryption Protocol ■ **3des**: 3DES-EDE Symmetric Encryption Protocol ■ **aes128**: 128- bit AES (Advanced Encryption Standard) ■ **aes192**:192- bit AES ■ **aes256**: 256-bit AES
<priv-pw>	Privacy password for the user. Passwords can have up to 32 characters; hyphens, underscores, and periods are acceptable.

Examples:

```
setup snmp remoteuser new "Jane Killick" auth md5 authpwd12$ priv aes245 privpwd12!
0000091E000000A1AC1512AC
setup snmp remoteuser new "Nonsecure user"
```

Delete a Remote User

To delete a remote user, use:

```
setup snmp remoteuser delete <username>
```

Example:

```
setup snmp remote user delete "Sean Wood "
```

View Remote User Settings

To view current remote user settings, issue the command:

```
setup snmp remoteuser show
```

Example output:

```
setup snmp remoteuser show

 RemoteUserName       AuthProt PrivProt RemoteEngineID                    Refs Status
 IT_remote            md5      aes256   0000091E000000A1AC1512AA           4 ok
 sys admin            none     none                                        3 ok
 Todd Gray            md5      des      0000091E000000A1AC1512A0           1 ok
```

> **Tip:** Since each individual SNMP setting change requires a full configuration reload, Symantec recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series
8.3.1	Command introduced

setup snmp show

Display the basic snmp configuration.

```
setup snmp show
```

The output is divided into SNMPv1/SNMPv2c and SNMPv3 settings. The first two sections of the output display settings for an SNMPv1 or SNMPv2 configuration. If you have configured SNMPv3 views, access groups, users and targets, the information for each table entry appears in the *ViewName*, *AccessGroupName*, *UserName, RemoteUserName* and *TargetName* sections.

The following is an example of the output of the **setup snmp show** command.

```
setup snmp show

SNMP config mode:       simple
SNMP look community:    lookpwd1
```

```
SNMP touch community:   touchpwd2
SNMP Trap destinations: 172.21.18.166 172.21.18.167

syslocation: Northwest Corner of Building 4
syscontact:  Jill Smith
sysname:     PKTR_9500_42
localSnmpID: 0000091E000000A1AC1512AA

ViewName             SubtreeOID                              Type              Refs Status
all_mib              1.3.6.1.6.3.1                           included          2       ok
isoAll               iso                                     included                  10      ok
isoNone              iso                                     excluded          1       ok
Traps                snmpTraps                               included          1       ok

AccessGroupName      Model Level        ReadViewName        WriteViewName       Refs Status
admin                usm   authPriv            all_mib             isoNone             1       ok
engineering          usm   authPriv            all_mib             Traps                       2       ok
test_1               usm   noAuthNoPriv        isoAll              isoAll                      1       ok
v1only               v1                        isoAll              isoAll                      1       ok
v2                   v2c                       isoAll              isoAll                      0       ok

UserName             GroupName            AuthProt                    PrivProt               Status
Amit                 engineering          md5                         des                    ok
IT                   engineering          md5                         des                    ok
Marcia               admin                md5                         des                    ok
RemoteUserName       AuthProt PrivProt RemoteEngineID                      Refs Status
IT_remote            md5      aes256   0000091E000000A1AC1512AA                              4 ok
sys admin            none     none
Todd Gray            md5      des      0000091E000000A1AC1512A0                              1 ok

TargetName       RemoteHost       RemoteUserName   ViewName         Type Ver Ml Status
System admin     10.10.14.55      IT_remote        isoAll           trap v3  um ok
Target_it        172.21.18.170    IT_remote        isoAll           trap v3  um ok
V1traps          172.21.18.160    public           isoAll           trap v1  v1 ok
```

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series

setup snmp syscontact

The name of the person managing the unit; *<string>* can be up to 256 characters long and must be enclosed in quotation marks if spaces are used. If the **default** parameter is specified, any existing syscontact value is cleared.

```
setup snmp syscontact <string>|default
```

Example:

```
setup snmp syscontact "Gail Jellison"
```

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series

setup snmp syslocation

Issue this command to specify the physical location of the PacketShaper (room, floor, building), where *<string>* can be up to 256 characters long and must be enclosed in quotation marks if spaces are used. If desired, specify the **default** parameter to clear the existing syslocation variable.

```
setup snmp syslocation <string>|default
```

Example:

```
setup snmp sysname "4th floor"
```

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series

setup snmp sysname

Issue this command to configure the PacketShaper's fully-qualified domain name for SNMPv1, where *<string>* can be up to 256 characters long and must be enclosed in quotation marks if spaces are used. If the **default** parameter is specified, the unit's IP address is used for the sysName variable.

```
setup snmp sysname <string>|default
```

Example:

```
setup snmp sysname example.com
```

If the PacketShaper is using its IP address as a default sysName and the unit's IP address changes, the sysName will update to the new IP address once the unit resets or an SNMP MIB walk is performed on the unit.

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series

setup snmp target

Create new SNMP targets or modify existing SNMP targets to determine where SNMPv3 notifications should be sent.

There are two different commands to modify target settings. Modify target settings on a PacketShaper in local mode or a top-level PolicyCenter configuration with the command **setup snmp target modify**. Use the command **setup snmp target override** on a PacketShaper in shared mode or a PolicyCenter child configuration to create a local copy of a SNMP target that overrides the inherited SNMP target.

```
setup snmp target modify|new|override <targetname> <targethost> <remoteuser> [port <port>]
[version v1|v2c|v3] [model v1|v2c|usm] [type trap|inform] [view <notifyView>][timeout
<seconds>] [retry <n>]
```

where

<targetname>	Name of the target you are creating or modifying. Target names can have up to 32 characters; hyphens, underscores, and periods are acceptable. If the target name includes a space, it must be enclosed within quotation marks, for example "target four." **Note:** A target name can be any text string, and does not have to be related to the target system or remote user name.
<targethost>	IP address of a remote IP host, in dotted-decimal format. If the new target will use both the **v3**protocol version *and* the *usm* security model, you must also specify a remote user.
<remoteuser>	To associate an existing remote user with this target, specify a user already defined by the command setup snmp remoteuser. To create a *new* remote user for this target with a localSnmpId and no authorization or privacy protection, specify the name for the new remote user. Note that remote users created with this command will not appear in the *Remote Users* table.
<port>	Port number on the remote host to which the notifications will be sent.
version *v1|v2|v3*	Specify **v1**, **v2** or **v3** to indicate which SNMP version of notifications the user will receive. The default SNMP version is **v3**.
model *v1|v2|usm*	Select the security model for this notification by specifying **v1** (for SNMPv1), **v2** (for SNMPv2c), or **usm** (for SNMPv3). The default security models for the different versions of SNMP are as follows: ■ SNMPv1: defaults to *v1* ■ SNMPv2c: defaults to *v2* ■ SNMPv3: defaults to *usm*
type trap|inform	Specify whether the user should receive trap notifications or just informs. If no parameter is specified, the default setting will be **trap**. **Note:** SNMPv1 supports trap notifications only.
view *<notifyview>*	To allow the remote user to receive all types of MIB notifications, specify the predefined view name **isoAll** for the *<notifyview>* parameter. To limit the user's access to a subset of available MIB notifications, enter the name of a user-defined view created with the **setup snmp view** command. If you do not specify a notify view, the group will apply the default **isoAll** setting.

timeout *<seconds>*	Maximum round trip time for communications between the PacketShaper and the SNMP target address, in seconds. Valid timeout values 1-60 , and the default value is **10**. If an inform message is sent to this address but a response is not received within this specified time frame, the PacketShaper will assume that there will be no response.
retry *<n>*	Number of times the PacketShaper should attempt to retransmit an inform message when it does not receive a response. Valid retry values are 1-10, and the default value is **3** retries.

Example:

```
setup snmp target new targ1 10.1.2.3 trapuser
```

Delete a Target

To delete a target, use:

```
setup snmp target delete <target>
```

Example:

```
setup snmp target delete "admin_target"
```

View SNMP Target Settings

To view current SNMP target settings, issue the command:

```
setup snmp target show
```

Example output:

```
TargetName       RemoteHost       RemoteUserName   ViewName         Type Ver  Ml Status
System admin     10.10.14.55      IT_remote        isoAll           trap v3  usm ok
Target_it        172.21.18.170    IT_remote        isoAll           trap v3  usm ok
V1traps          172.21.18.160    public           isoAll           trap v1   v1 ok
```

> **Tip:** Since each individual SNMP setting change requires a full configuration reload, Symantec recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series
8.3.1	Command introduced

setup snmp user

Each SNMP user entry defines a user (login) name, an association with an existing access group, and authentication and privacy keys that a management system can use to access the PacketShaper. This user name is not related to any other user names such as those defined for RADIUS or PolicyCenter access.

There are two different commands to modify user settings. Modify user settings on a PacketShaper in local mode or a top-level PolicyCenter configuration with the command **setup snmp user modify**. Use the command **setup snmp user override** on a PacketShaper in shared mode or a PolicyCenter child configuration to create a local copy of a SNMP user that overrides the inherited SNMP user.

> **Note:** If you have not yet defined access groups for your SNMP users, use the CLI command setup snmp accessgroup to create one or more access groups before you add users to these groups.

```
setup snmp user new|modify|override <username> <groupname> [auth {md5|sha} <auth-pw>] [priv {des|3des|aes128|aes192|aes256} <priv-pw>]
```

where

<username>	Name of the user you are creating or modifying. A user name can have up to 32 characters; hyphens, underscores, and periods are acceptable. If the name includes a space, it must be enclosed within quotation marks, for example "Jane Doe." Each SNMP user name must be unique.
<groupname>	User's access group
auth {md5|sha}	If the user's access group uses the *usm* (SNMPv3) security model with the *authNoPriv* or *authpriv* security levels, Specify either the **MD5** or **SHA** authorization protocol and enter an authentication password for the user. If the user's access group uses the *v1* (SNMPv1) or *v2* (SNMPv2c) security model or an *noAuthNoPriv* security level, this parameter is not required.
<auth-pw>	Authorization password for the user. Passwords can have up to 32 characters; hyphens, underscores, and periods are acceptable.
priv {des|3des|aes128|aes192|aes256}	Specify one of the following privacy protection protocols only if the user's access group uses the*authpriv* security level. Otherwise, this parameter is not required. ■ **des**: CBC-DES Symmetric Encryption Protocol ■ **3des**: 3DES-EDE Symmetric Encryption Protocol ■ **aes128**: 128- bit AES (Advanced Encryption Standard) ■ **aes192**:192- bit AES ■ **aes256**: 256-bit AES
<priv-pw>	Privacy password for the user. Passwords can have up to 32 characters; hyphens, underscores, and periods are acceptable.

Examples:

```
setup snmp user new "Kim Johnson" snmpv3Eng auth md5 authpwd123 priv aes245 privpwd123

setup snmp user new "v1_user" snmpv1Group

setup snmp user modify "Kim Johnson" snmpv3Eng auth md5 new_pwd1 priv aes245 new_pwd2
```

> **Tip:** Since each individual SNMP setting change requires a full configuration reload, Symantec recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

Delete a User

To delete a user, use:

```
setup snmp user delete <username>
```

Example:

```
setup snmp user delete "Ken Traum"
```

View Users' Group and Security Settings

To view current SNMP user settings, issue the command:

```
setup snmp user show
```

Example output:

```
 setup snmp user show
   UserName             GroupName            AuthProt PrivProt Status
   Amit                 engineering          md5      des      ok
   Example_v1           v1only               none     none     ok
   IT                   engineering          md5      des      ok
   Jane Doe             engineering          md5      des      ok
   Kim Johnson          admin                md5      des      ok
   Tom Jones            authnopriv           md5      none     ok
   VP_Marcia            admin                md5      des      ok
   Wendy Ho             engineering          md5      des      ok
```

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series
8.3.1	Command introduced

setup snmp view

An SNMP view filters objects from the entire MIB and defines a subset of MIB objects. Every SNMP access group has views for read and write access which either allow or limit that group's access to MIB objects. There are two predefined views; *isoAll* and *isoNone*. The *isoAll* view gives a group access to all MIB information, and the *isoNone* view blocks all access.

If you want your SNMP groups to have either complete access or no access to all MIB information, your groups only need to use the built-in *isoAll* or *isoNone* views. If, however, you want a group to access just a subset of MIB information, you will have to create a new view that describes those MIB object identifiers (OIDs) that should be included or excluded.

There are two different commands to modify view settings. Modify view settings on a PacketShaper in local mode or a top-level PolicyCenter configuration with the command **setup snmp view modify**. Use the command **setup snmp view override** on a PacketShaper in shared mode or a PolicyCenter child configuration to create a local copy of a SNMP view that overrides the inherited SNMP view.

```
setup snmp view add|modify|new|override <viewName> <OID> [included|excluded]
```

where

<viewname>	Name of the view you are creating or modifying. A view name can be up to 32 characters; hyphens, underscores, and periods are acceptable. If the name includes a space, it must be enclosed within quotation marks, for example "admin view." Each SNMP view name must be unique.
<OID>	The *<OID>* parameter may be an OID name, number, or an initial OID name and a number, e.g., *packeteerMibs*, *1.3.6.1.4.1.2334.2* or *packeteerMibs.1.4* This parameter also supports the use of asterisks as wildcards for OID numbers, for example, *interfaces.*.*.1*
included|excluded	Select **included** if the OID subtree should be included in this view, or **excluded** if it is explicitly not accessible. If you do not specify an include or exclude parameter, the OID will automatically be included.

The `setup snmp view new` command only allows you to specify a single OID. To include or exclude an additional OID in the view, use the command `setup snmp view add`.

Examples:

```
setup snmp view new sysadmin 1.3.6.1.6.3.18 exclude
setup snmp view add sysadmin 1.3.6.1.6.3.15.1.2.2 exclude
```

When you modify snmp view settings, all OIDs defined for that view are removed and replaced with the one OID specified in the **setup snmp view modify** or **setup snmp view overide** command. To add additional OIDs to the modified view, use **setup snmp view add**.

Examples:

```
setup snmp view modify sysadmin 1.3.6.1.6.3.19 exclude
setup snmp view add sysadmin 1.3.6.1.6.3.15.1.2.2 exclude
```

Tip: Since each individual SNMP setting change requires a full configuration reload, Symantec recommends that you temporarily disable SNMP (or wait to enable it) before modifying multiple SNMP settings. With this approach, the configuration only needs to reload once (upon enabling SNMP) which avoids significant overhead.

Delete a View

To delete a view, use:

```
setup snmp view delete <viewName>
```

Example:

```
setup snmp vew delete IT_view
```

Display View Settings

To display current SNMP view settings, issue the command:

```
setup snmp view show
```

Example output is shown below. The letter **B** before the *isoAll* and *isoNone* view names indicate that these views are built-in, and cannot be modified or deleted.

```
 setup snmp view show

  ViewName             SubtreeOID                Type      Refs Status
  adminview            1.3.6.1.6.3.1            included    8 ok
  iface_view           1.3.6.1.2.1.2            included   12 ok
B isoAll               iso                      included   24 ok
B isoNone              iso                      excluded    2 ok
  sysadmin             1.3.6.1.6.3.15.1.2.2     excluded    3 ok
                       1.3.6.1.6.3.19           excluded
  Traps                snmpTraps                included    5 ok
```

Command Change History

Release	Modification
11.2.1	SNMP support added for PS S-series
8.3.1	Command introduced

setup sntp

Set or display the Simple Network Time Protocol (SNTP) configuration. SNTP is used to synchronize the time in PacketShaper to a server configured to propagate highly accurate time information through the Internet.

`setup sntp on|off|servers {<primary> [<secondary>]|none}|poll|reset|sync`

Enter a standard dotted-decimal IP address.

setup ssl install

Import an SSL certificate in the PEM encoding format, PacketShaper has its own self-signed certificate, but if you prefer, you can generate your own PEM file using OpenSSL library, and install the certificate with the **`setup ssl install`** command.

```
setup ssl install <pem file>
```

The PEM file must be located in the public/3.0/ssl folder on your PacketShaper before you issue the command.

> **Tip:** Use the File Browser to upload the PEM file to the public/3.0/ssl folder.

Additional Information:

- Each PacketShaper has a unique self-signed certificate, identified by its serial number.
- PacketShaper will validate the PEM file during its installation and will alert you if an error or corruption is discovered.
- For security purposes, be sure to delete the PEM file from other locations after it is installed on PacketShaper.

Command Change History

Release	Modification
11.5.3	Command introduced.

setup standby

Enable/disable a PacketShaper for high availability. PacketShaper S-Series appliances can be installed in redundant network paths to provide link redundancy in case a component (PacketShaper, router, or switch) fails. All PacketShaper units in a configured high availability cluster are active and manage traffic on their own links, but they are simultaneously and continuously monitoring the health and status of its partners to make sure a unit has not gone down.

```
setup standby on|off|show
```

Where:

on	Enables the unit for high availability
off	Disables the unit for high availability
show	Displays the standby status (on or off), and lists the unit's standby partners.

To enable high availability:

```
setup standby on
```

To disable high availability:

```
setup standby off
```

To show the status of the standby configuration:

```
setup standby show

Standby Status: ON

Standby Configuration:
ID  Serial No       IP Address      State Down   Uptime   Sent Bytes   Recv Bytes
0 03xxxxxxxxx      10.9.64.155    running    0  3135819            0   1797149974
```

The following details are provided for each partner: serial number, IP address, state, up/down status (0=up, 1=down), uptime in seconds, number of bytes sent and received since standby has been running. The State field is useful when setting up or troubleshooting the feature. Possible states include:

- **init**: initial state
- **conn**: connecting
- **negotiate**: negotiating after connection is successful
- **sync**: synchronizing configuration and other data
- **running**: the running state
- **inactive**: disabled

Command Change History

Release	Modification
11.6.1	Command introduced

setup standby partner

Add or delete partners in a high availability cluster.

```
setup standby partner add|del <partner_ip>
```

To add a standby partner:

```
setup standby partner add 10.9.64.12
```

To delete a partner:

```
setup standby partner del 10.9.64.12
```

Command Change History

Release	Modification
11.6.1	Command introduced

setup syslog add

Add a Syslog server. The logging feature gives administrators a way to centrally log and analyze user events and system warning messages. For example, if you are using RADIUS authentication, each failed login attempt will be sent to the defined Syslog server.

Adaptive response action files and user events can be configured to send messages to a Syslog server. For example, when you register an event, you will be asked if you want to send events to Syslog; you can define and register an event that sends a message to a Syslog server when retransmissions rise to 30 percent of your network activity.

You can add up to four servers.

```
setup syslog add host:<ipaddress> [output:<facility>,<level>] [port:<portnum>] [datetime]
```

host: *<ipaddress>*	The Syslog server IP address for example, *host:10.7.38.100*
	The facility and severity level for example, *output:local1,6*
output: *<facility>* ,*<level>*	Up to three outputs can be specified. The default facility is local4 and the default level is 7. PacketShaper user events are at severity level 6; if you want to capture them with Syslog, you must set the level to 6 or 7. See Facility Types and Severity Levels for lists of the valid facility types and levels.
port: *<portnum>*	The port number of the Syslog server; if the port isn t specified, port 514 is used
datetime	Include the date and time in the message; the date and time are not included unless you specify the **datetime** parameter

For example:

```
setup syslog add host:10.7.38.100 output:local1,3 datetime
```

If you need to modify any of the settings later, you need to remove the server and then add it again (see setup syslog remove).

Messages are not sent until you enable the logging feature. See setup syslog state on. If you want a PacketShaper event to be recorded in a Syslog, you need to specify this option when registering the event (see event register).

Facility Types

You can enter the keyword or value specified in the following table.

Description	Keyword	Value
Kernel	kern	0
User Processes	user	1
Electronic Mail	mail	2
Background System Processes	sysd	3
Authorization	auth	4
System Logging	sysl	5
Printing	lpr	6
Usenet News	news	7

Description	Keyword	Value
Unix-to-Unix Copy Program	uucp	8
Clock Daemon	clkd	9
Security	sec2	10
NTP Subsystem	ntp	12
Log Audit	audit	13
Log Alert	alert	14
Clock Daemon	clkd2	15
For Local Use	local0 local7	16-23

Severity Levels

You can enter the keyword or value specified in the following table. Set the level to specify which messages to suppress to the Syslog server. For example, setting the severity level to 3 allows messages with levels 0 3 and suppresses messages with levels 4 7. If you don't specify a severity level, 7 is used. With the default severity level, messages of all levels will get sent to the Syslog server.

Description	Keyword	Value
System unusable	emerg	0
Take immediate action	alert	1
Critical condition	crit	2
Error message	err	3
Warning message	warn	4
Normal but significant condition	notice	5
Informational	info	6
Debug message	debug	7

At the "warn" level, Symantec sends the following types of messages to the Syslog server:

- Login failed
- Hard drive status
- Measurement Engine status

User events that are configured to send a syslog message when a threshold is crossed are sent at the info severity level (6). See event register for more information on configuring an event to send a syslog message.

Adaptive response action files that include the send syslog command can designate the severity level at which the message is sent to the Syslog server; any level can be specified.

setup syslog rate

Set the maximum number of syslog messages that will be sent per second.

```
setup syslog rate <number>
```

The default rate is 20 messages per second and the valid range is 1-200. You may want to increase the rate if you are experiencing a problem with your unit.

setup syslog remove

Remove a syslog server. If you need to modify the settings of a server you have added, you will need to remove the server first.

```
setup syslog remove <ipaddress>
```

setup syslog show

Display the settings for currently defined Syslog servers.

```
setup syslog show [<ipaddress>]
```

If no *<ipaddress>* is specified, the setup of all Syslog servers is displayed. For example:

```
setup syslog show
```

```
        Status: On
 Max Rate: 35
 Total Sent: 5
 Total Lost: 0

Server Addr                Facility             Level
-----------------------------------
10.7.38.200                local4, 20           warn, 4
10.7.38.100                local4, 20           warn, 4
```

If you specify an *<ipaddress>*, the settings for a single Syslog server are displayed. For example:

```
setup syslog show 10.7.38.200
```

```
    Server Addr: 10.7.38.100
 UDP Port: 514
 DateTime Option: Not Enabled

-----------------------------------
Facility                           Level
-----------------------------------
local4, 20                         warn, 4
```

Message Format

When viewing the messages at the Syslog server, you will see the format of a Syslog message is as follows:

```
ReceiveDateTime address SendDateTime module-severity-MNEMONIC: description
```

ReceiveDateTime	The date and time the message was received by the Syslog server (may not be included, depending on the setup of the Syslog server)
address	The PacketShaper unit s IP address
SendDateTime	The date and time the message was sent to the Syslog server (if the datetime parameter was specified when defining the syslog server)
module	A four-byte string that identifies the type of message. For example, USRE is a user event and SYSW is a system warning.
severity	A single digit code (0 7) that reflects the severity of the condition; see Severity Levels
MNEMONIC	A code that uniquely identifies the error message for example, BAD_WR (bad write) or INSERT_F (insert into a list fails)
description	A text string describing the condition

Example message:

```
Aug 6 17:06:27 10.7.38.5 SYSW-4-LOG_WARN: Hard drive is down.
```

Or, if the datetime parameter was specified:

```
Aug 6 17:07:25 10.7.38.5 Mon Aug 6 17:05:01 2001 BST (London) SYSW-4-LOG_WARN: Hard drive
is down.
```

setup syslog state

Enable or disable the logging feature so that messages will be sent to the defined syslog server(s). If this command is issued from the command-line interface of an individual PacketShaper or a PolicyCenter sharable configuration, it will enable or disable the syslog servers defined for the selected configuration or unit.

```
setup syslog state on|off|default
```

Select the **default** option to set the logging feature on a PacketShaper to its default *off* state. To check whether the logging feature is on or off, use the setup syslog show command.

> **Note:** If the **setup syslog state default** command is issued from PolicyCenter for a *child configuration*, the selected configuration will discard its existing syslog state and inherit its syslog on/off setting from its parent configuration. If the **setup syslog state default** command is issued from PolicyCenter for a *root-level configuration*, the syslog state for that configuration will be returned to the default *off* setting.

setup tacacs acct

Set up or change the configuration of the TACACS+ accounting service records. This feature allows you to have an audit trail for user logins.

To define the TACACS+ accounting service, use:

```
setup tacacs acct primary|secondary {<host> <shared_secret> [<port>]}|delete|override
```

primary|secondary	Enter the literal primary or secondary to indicate which server you are defining. (The secondary server is used when the primary server isn t accessible.)
<host>	The IP address or DNS name of the TACACS+ accounting server
<shared_secret>	The designated secret for the server; quotes are not required
[*<port>*]	The port number to access the server; if omitted, the default port *49* is used.
delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	For a PolicyCenter child configuration, create local TACACS+ settings that override the settings it inherits from its parent configuration. Remove these override settings at any time with the command **setup tacacs acct primary|secondary delete**.

To turn the service on or off, or to return the service to its default **off** value, use:

```
setup tacacs acct on|off|default
```

Example:

```
setup tacacs acct primary 10.10.10.10 P4assw0rd1

setup tacacs acct secondary 10.10.20.10 Paa55w0rd2

setup tacacs acct on
```

This example defines a primary accounting server at 10.10.10.10 which has a shared secret of P4ssw0rd1, as well as a secondary server at 10.10.20.10. The third command line enables the TACACS+ accounting service. Once this service is configured and enabled, PacketShaper will send a PW_STATUS_START accounting message to the accounting server when a user logs in and a PW_STATUS_STOP message when a user logs off or is disconnected.

Command Change History

Release	Modification
11.5.1	Command introduced

setup tacacs auth

Set up or change the configuration of the TACACS+ authentication service. TACACS+ authentication is an optional method for users to log into the PacketShaper browser interface, command-line interface, or customer portal or when FTPing to the unit. Using third-party TACACS+ servers enables you to have central configuration of user accounts.

```
setup tacacs auth primary|secondary {<host> <shared_secret> [<port>]}|delete|override
```

primary|secondary	Enter the literal primary or secondary to indicate which server you are defining. (Note: The TACACS+ client uses the secondary server when the primary server isn't accessible or authentication failed.)
<host>	The IP address or DNS name of the TACACS+ authentication server
<shared_secret>	The designated secret for the server; quotes are not required

[*<port>*]	The port number to access the server; if omitted, the default port *49* is used
delete	Deletes the configuration of the primary or secondary server (whichever is specified)
override	For a PolicyCenter child configuration, create local TACACS+ settings that override settings inherited from its parent configuration. Remove these override settings at any time with the command **setup tacacs primary|secondary delete**.

To turn the service on or off, or to return the service to its default **off** value, use:

```
setup tacacs auth on|off|default
```

Example:

```
setup tacacs auth primary 10.10.10.10 CupServ44

setup tacacs auth on
```

This example first defines a primary authentication server at 10.10.10.10 which has a shared secret of CupServ44. The second command line enables TACACS+ authentication service. Once this is configured and enabled, PacketShaper will prompt users for user name and password when they log into PacketShaper.

setup tacacs clear

Clear the TACACS accounting request queue and reset the count of dropped requests. This command will clear the banner message that appears after TACACS accounting requests have been dropped.

```
setup tacacs clear
```

Command Change History

Release	Modification
11.5.1	Command introduced

setup tacacs method

Select the TACACS+ authentication method:

- ***ASCII*** *(American Standard Code for Information Interchange)*: With ASCII, the username and password are transmitted in clear, unencrypted text.

- ***PAP*** *(Password Authentication Protocol)*: With PAP, the username and password are transmitted in clear, unencrypted text. ASCII or PAP authentication is required for TACACS+ configurations that require access to clear text passwords (for example, when passwords are stored and maintained in a database external to the TACACS+ server)

- ***CHAP*** *(Challenge Handshake Authentication Protocol):* In other environments, CHAP may be preferred for greater security. The TACACS server sends a challenge that consists of a session ID and an arbitrary challenge string, and the username and password are encrypted before they are sent back to the server.

- ***MS-CHAP*** *(Microsoft Challenge Handshake Authentication Protocol):* This protocol is very similar to CHAP, but with MS-CHAP authentication, the TACACS+ server can store an encrypted version of a user password to validate the challenge response. Standard CHAP authentication requires that the server stores unencrypted passwords.

  Note: MS-CHAP v1 and v2 are supported. PacketShaper attempts authentication with MS-CHAP v2 first. If the remote server doesn't support v2 or if authentication is denied, PacketShaper re-attempts authentication with MS-CHAP v1.

```
setup tacacs method ascii|pap|chap|mschap|default
```

The default authentication method is **ascii**.

Command Change History

Release	Modification
11.5.1	Command introduced

setup tacacs show

Display the current TACACS+ settings. Use this command to verify that TACACS+ authentication and accounting are enabled, to see the timeout setting, and to view configuration settings on each of the TACACS+ servers.

```
setup tacacs show

TACACS Setup values:
     Tacacs Method   : ASCII
     Authentication  : on
     Accounting      : off
     Timeout         : 10

  TACACS Service records:
     Type  Host                                    Port Secret
     auth1 192.21.18.190                             49 test
```

Command Change History

Release	Modification
11.5.1	Command introduced

setup tacacs timeout

Set the amount of time for PacketShaper to wait for a response from the TACACS+ server. If the server doesn't send a reply within the timeout period, the PacketShaper will disconnect and the authorization attempt will fail. The default timeout period is 10 seconds.

```
setup tacacs timeout <seconds>|default
```

where *<seconds>* is a value between 1 and 60 seconds.

In the example below, the timeout interval is 25 seconds; this interval applies to any configured TACACS+ server.

```
setup tacacs timeout 25
```

To return to the default timeout interval, use:

```
setup tacacs timeout default
```

Command Change History

Release	Modification
11.5.1	Command introduced

setup timezone

When you configure a time zone, PacketShaper can change its local time automatically at the start and end of daylight savings time. It also can retrieve time updates from time servers.

```
setup timezone [<name>|custom <timezone_spec>]
```

Each time zone has a unique name usually the name of the best-known city in that zone. The default time zone is Los Angeles, CA. To display the valid time zones, use `setup timezone help`.

<timezone_spec> is a string defined by POSIX.1 as:

<std><offset>[*<dst>*[*<offset>*],*<date>*[/*<time>*],*<date>*[/*<time>*]]

Where:

<std> and *<dst>*	3 or more characters specifying the standard and daylight saving time (DST) zone names
<offset>	[-]hh:[mm[:ss]] specifies the offset west of UTC. The default DST offset is one hour ahead of standard time
<date>[/*<time>*]	Specifies the beginning and end of DST. If this is absent, the system applies US DST rules (first Sunday of April at 2:00 AM to last Sunday of October at 2:00 AM)
<time>	hh:[mm[:ss]] with a default of 02:00
<date>	One of the following forms: Jn (1<=n<=365): origin-1 day number, not counting February 29 n (0<=n<=365): origin-0 day number, counting February 29, if present Mm.n.d (0[Sunday]<=d<=6[Saturday], 1<=n<=5, 1<=m<=12): for the dth day of week n of month m of the year, where week 1 is the first week in which day d appears, and 5 stands for the last week in which day d appears (which may be either the 4th or 5th week)

For example, you could configure a time zone for Cairo, Egypt with the command:

```
setup timezone custom EET-2EEST,M4.5.5/01:00,M9.5.5/03:00
```

```
setup timezone
```

Current time zone:
Time zone name: Custom
Time zone desc: Custom time spec in POSIX format
Time zone spec: EET-2EEST,M4.5.5/01:00,M9.5.5/03:00
Time zone offset: GMT+02:00
DST offset: 60 minutes
DST starts: Last Friday of April at 01:00 AM
DST ends: Last Friday of September at 03:00 AM

In this example, the standard time, known as EET, is two hours ahead of GMT and daylight savings time, known as EEST, is the default 60 minutes ahead of EET. Rather than using US default rules, EEST begins on the last Friday of April at 1:00 AM and ends on the last Friday of September at 3:00 AM.

> **Note:** You should always do a system reset immediately after changing the time zone so that the underlying time-sensitive scheduled operations of the PacketShaper can be correctly initialized.

setup variable

Change a default variable setting.

```
setup variable [<variable> <value>|default] | [-reset|-nd]
```

where *<variable>* is one of the variables listed below and *<value>* is the value you want to set the variable to. The default, minimum, and maximum values for each *<variable>* are listed in the table.

> **Note:** After changing a variable's setting, many variables require that you reset the unit in order for the change to take effect.

To reset all system variables to their defaults, use the **setup variable -reset** command. To reset a specific variable to its default, use the **setup variable *<variable>* default** command. To see a list of all variables that have non-default settings, use the **setup variable -nd** command.

Variable/ Description	Default Value	Min. Value	Max. Value
autoCreateSameSide When this variable is enabled, the SameSide class is created automatically. When disabled, the SameSide class will not be auto-created. You may want to disable this variable if traffic is being misclassified into the SameSide class.	1 (on)	0 (off)	1 (on)
Bridge PassThru With Bridge PassThru enabled, the PacketShaper forwards packets that have a source and destination MAC address on the same side of the unit. When Bridge PassThru is disabled and traffic shaping is enabled, the PacketShaper drops packets that have source and destination MAC addresses on the same side.	1 (on)	0 (off)	1 (on)
DiffservClassSortPref Controls the sort order of the traffic tree, with respect to Diffserv classes (those with DSCP marks). Three settings are available: **0** Diffserv classes are sorted below IP-address-based classes, but above port-based classes (the default). **1** Diffserv classes are sorted above IP-address-based classes **2** Legacy sort order (Diffserv classes are sorted after IP-address-based classes, port-based classes, and auto-discovered classes) **Note:** The new sort order doesn't take effect until the unit is rebooted.	0	0	2
discoveryThresholdDynamicPort The number of new connections of an identifiable service to a port greater than 1024 that must be identified within a one-minute timeframe before PacketShaper creates a class	2	1	1000000
discoveryThresholdNonIP The number of new non-IP connections of a given type that must be identified within a one-minute timeframe before PacketShaper creates a class	2	1	1000000
discoveryThresholdNormal The number of new connections of an identifiable service to a port less than or equal to 1024 that must be identified within a one-minute timeframe before PacketShaper creates a class	1	1	1000000
discoveryThresholdPort The number of new connections to a particular port within a one-minute timeframe before PacketShaper creates a Port_#### class in the DiscoveredPorts folder It may be necessary to increase this value on Internet link deployments to prevent excessive number of DiscoveredPorts classes being created. If you don t want any Port_ #### classes discovered, set this variable to its maximum value.	100	1	1000000
discoveryThresholdUrlCategories The number of new flows belonging to a particular URL category that must be identified within a one-minute time frame before PacketShaper creates a class for the category	1	1	1000000
dynPtnActiveReuseSeconds The number of seconds a dynamic partition will be retained after an established flow has sent packets. (Flow is not terminated/completed but hasn't had any recent activity.) **Note:** If no other user needs a dynamic partition, the partition will be retained indefinitely.	300 (5 min)	10	7200 (2 hrs)
dynPtnIdleReuseSeconds The number of seconds a dynamic partition will be retained after all flows in the partition have been terminated/completed. **Note:** If no other user needs a dynamic partition, the partition will be retained indefinitely.	30	10	7200 (2 hrs)

Variable/ Description	Default Value	Min. Value	Max. Value
dynPtnSequestrationCount The number of partitions reserved for static partitions; all other partitions can be used for dynamic or static partitions	3	0	99
enableCongestion Enable/disable the calculation of packet exchange time. When this variable is disabled, the Pkt Exch column on the Monitor Traffic page will not appear, RTM will not be available, and the packet exchange time and RTM measurement variables will always have a value of 0. After disabling the enableCongestion variable, you should reset the unit.	1 (on)	0 (off)	1 (on)
enableJumboFrame Enable/disable jumbo frame support. When enabled, the PacketShaper NIC driver supports jumbo frames (9000-byte MTU). When the variable is disabled, the NIC driver supports only standard frames (1500-byte MTU) and will drop packets that exceed this size. This feature has the following minimum memory requirements: PS-S200: 16GB PS-S400: 32GB PS-S500: 64GB **Note:** When the jumbo frame variable is enabled, you must also modify the inbound and outbound maximum segment size. Change the **tcpMssInbound** and **tcpMssOutbound** system variables from their default settings of 1460 bytes to 8960 bytes (equivalent to the payload size of jumbo frames).	0 (off)	0 (off)	1 (on)
enableLatency Enable/disable the calculation of VoIP metrics. When this variable is enabled, PacketShaper collects data that measure packet loss, jitter, and latency for VoIP flows. **Additional Information** ■ VoIP metrics can only be measured between PacketShapers with the VoIP metrics feature enabled. ■ The VoIP metrics feature can measure traffic only from VoIP applications whose data is classified as RTP-I. For instance, latency metrics are not provided for DialPad, iChat, Vonage, and Skype.	0 (off)	0 (off)	1 (on)
enableSTUNclassification Enable/disable classification of the STUN (Session Traversal Utilities for NAT) protocol. Audio/video flows are first classified as STUN and then as RTP/RTCP. When enabled, PacketShaper will auto-discover the STUN class. If your network has a lot of STUN traffic, you may want to disable STUN classification to improve performance. **Note:** RTP/RTCP will still be classified even when STUN classification is disabled.	1 (on)	0 (off)	1 (on)

Variable/ Description	Default Value	Min. Value	Max. Value
enableVoIPUseragentAutoDiscovery RTP auto-discovery is based on the VoIP user-agent attribute when this variable is enabled. When this variable is disabled (as it is by default), RTP auto-discovery is based on the RTP-I encoding attribute. When this variable is enabled and RTP-I is auto-discovered, it will auto-discover child classes based on VoIP user agent traffic (such as RTP-I-Motorola_VT1000 and RTP-I-Google_Talk).	0 (off)	0 (off)	1 (on)
enableWinnyClassification Enable/disable classification of the Winny service. For optimal performance, enable only when management of Winny traffic is required. **Note:** The Winny peer-to-peer application is used primarily in Japan.	0 (off)	0 (off)	1 (on)
flowRecordsIntermediateTimeout Number of milliseconds between generation and sending of intermediate flow detail records when traffic is present	1500	1000	36000
flowRecordsPktr0Timeout Number of seconds between generation and sending of Packeteer-0 flow records.	3600	10	5000
flowRecordsPktrPTimeout Number of seconds between generation and sending of Packeteer-P flow records.	60	10	5000
flowRecordsResetCounters Controls whether or not the counter fields in FDR packets are reset with each intermediate FDR sent **Note:** This variable only affects Packeteer-1 and Packeteer-2 format FDRs: counter fields are always reset in the NetFlow-5 format.	1 (on)	0 (off)	1 (on)
flowRecordsSendIntermediate Enable/disable the intermediate flow detail records feature. When this variable is enabled, PacketShaper emits intermediate FDRs at the interval specified by the **flowRecordsIntermediateTimeout** variable. **Note:** Enable the intermediate flow detail records feature only when using a suitably-instrumented collector, such as Cisco-based Netflow-5 collectors. IntelligenceCenter does not support intermediate FDRs.	0 (off)	0 (off)	1 (on)
flowRecordsSendPktrP Enable/disable emission of Packeteer-P packets to Packeteer-1 and Packeteer-2 flow detail record collectors. Packeteer-P packets contain statistics that are not related to particular flows, but rather provide information about utilization on the PacketShaper at the time flows are recorded. If this variable is enabled, Packeteer-P records are sent after each UDP flow record packet is sent to Packeteer-1 or Packeteer-2 collectors (not more than once per minute).	0 (off)	0 (off)	1 (on)

Variable/ Description	Default Value	Min. Value	Max. Value
flowRecordsSendPktr0 Enable/disable emission of Packeteer-0 packets to Packeteer-1 and Packeteer-2 flow detail record collectors. Packeteer-0 packets are mapping messages that allow collectors to decipher PacketShaper-related information in the FDRs they receive. For example, in the FDR s ClassID field, a value identifies the traffic class. In order for the collector to understand what class is actually associated with the ID, it uses the class map a list that contains each traffic class on the unit along with the identifying number assigned to each class. If this variable is enabled, Packeteer-0 mappings are sent out approximately once each hour. Note that this variable needs to be enabled only if the collector does not know this information through other means.	0 (off)	0 (off)	1 (on)
graphTimeoutSeconds The maximum number of seconds a graph can take to generate in the browser interface; if the graph takes longer to generate than this value, a system timeout error message will appear. **Note:** Increasing this setting can make the browser interface appear to "freeze" while PacketShaper is generating some of the more complex graphs. Sometimes the browser will not display the page until all of the graphs are generated.	60	1	600 (10 min)
hostTspecCacheInside Enable/disable caching of IP address-based classes on the inside. Change this setting to outside (0) to increase performance of classification if the majority of IP addresses in manually created classes are on the outside, rather than the inside. To disable the caching of inside IP address-based classes, use the **setup variable hostTspecCacheInside 0** command. After you reset the PacketShaper, IP address-based classes will be cacheable on the outside. To re-enable caching for inside classes, use the **setup variable hostTspecCacheInside 1** command.	1 (inside)	0 (outside)	1 (inside)
httpStealth503 Control the display of the '503 - Service unavailable' server error message when a connection is refused because of admission control (such as a never-admit policy). **0** The '503 - Service unavailable' message will be customized with the text 'This message is sent by Symantec PacketShaper.' **1** The PacketShaper text is not displayed with the '503 - Service unavailable' message. **2** PacketShaper performs a TCP reset and drops the HTTP request; the error message will likely be 'The attempt to load *http://...* failed.'	0	0	2
ipUserCacheNegativeTTL The number of seconds an IP will be stored in the PacketShaper cache, when the IP lookup does not result in a user.	1800	300	86400
ipUserCachePositiveTTL The number of seconds an IP-user name mapping will be stored in the PacketShaper cache. By default, the user mappings are stored for one hour. Because querying the cache is faster than querying the BCAAA server, you can accelerate user name look ups by increasing the cache timeout. However, the tradeoff is that stale mappings could cause incorrect user name identification.	3600	300	86400
latencyProbeDiscard Allows the PacketShaper to be configured to discard VoIP latency probes after responding. If VoIP devices located on the Inside of the PacketShaper are sensitive to VoIP latency probes, enabling this variable will prevent potential VoIP call drops.	0	0	1

Variable/ Description	Default Value	Min. Value	Max. Value
LFNSupport When enabled, this setting improves performance on Long Fat Networks (LFN) which require larger TCP window sizes. An LFN is a long distance network with large bandwidth and long delay; for example, high-capacity satellite channels are LFNs.	0 (off)	0 (off)	1 (on)
linkOverheadBytes Number of bytes that are added to each packet to account for WAN protocol header overhead	0	0	256
linkOverheadPpt Number of parts per thousand* by which packet sizes are increased to account for link overhead. This adjustment is useful for links that do bit stuffing. (*Bit stuffing* is the practice of adding bits to a stream of data. Bit stuffing is required by many network and communications protocols, for example to prevent data from being interpreted as control information.) * to be more precise, it's actually parts per 1024	0	0	1024
mirrorLinks Enable/disable link state mirroring. With link state mirroring, PacketShaper will bring down the second port of a NIC pair if the first goes down. This feature allows each PacketShaper to sit between a WAN router and a switch without blocking detection of switch outages by the router.	0 (off)	0 (off)	1 (on)
padVlanPkt Prevent PacketShaper from stripping out zero padding from 68- byte frames with VLAN tags. When the variable is disabled, PacketShaper behaves as it previously did when it received a 68-byte VLAN-tagged packet with zero padding: it strips out the zero padding in the packet. Because this packet modification can cause the packet to be dropped in some deployments, enabling the system variable can prevent this behavior.	0 (off)	0 (off)	1 (on)
PolicyFlowLimitForAllClasses Enables/disables the policy flow limit feature. When enabled, PacketShaper will enforce all policy flow limits that have been set on traffic classes. When disabled, all policy flow limits will be ignored. Disabled is the appropriate setting for PacketShapers deployed in proxy or NAT environments. For additional information, see policy flowlimit.	0 (off)	0 (off)	1 (on)
probeIntervalSeconds Number of seconds between the issuance of VoIP latency probes that measure VoIP metrics, enabled by the **enableLatency** variable.	5	1	60
rtoInboundClampMsecs Number of milliseconds delay for clamping early retransmission timeout on Inbound packets. Puts a maximum on retransmit time.	1600	0 (disable)	3000 (3 sec)
rtoOutboundClampMsecs Number of milliseconds delay for clamping early retransmission timeout on Outbound packets.	1600	0 (disable)	3000 (3 sec)
syntheticReadTimeoutSeconds Number of seconds after which a synthetic transaction will end when the response received is incomplete	5	1	1000
syntheticWriteTimeoutSeconds Number of seconds after which a synthetic transaction will be canceled if the server fails to respond to a request	60	10	5000

Variable/ Description	Default Value	Min. Value	Max. Value
tcpClipInitialWindow When tcpClipInitialWindow is enabled, the PacketShaper will always reduce the initial TCP window size to 1x MSS (maximum segment size). When this variable is disabled, new flows will ramp up faster but enforcement of small rate policies and/or partitions may not work at the beginning of flows.	1 (on)	0 (off)	1 (on)
tcpMssInbound Maximum segment size of TCP packets on Inbound flows. This setting can help avoid packet fragmentation when using VPN and not being able to support 1500-byte packets (the default size) through the VPN tunnel. **Note:** When the **enableJumboFrame** variable is enabled, change **tcpMssInbound** to 8960 bytes (equivalent to the payload size of jumbo frames).	1460 bytes	0	65535
tcpMssOutbound Maximum segment size of TCP packets on Outbound flows. **Note:** When the **enableJumboFrame** variable is enabled, change **tcpMssOutbound** to 8960 bytes (equivalent to the payload size of jumbo frames).	1460 bytes	0	65535
tcpSmallMssLinkSpeed Link speeds slower than this value will force the use of smaller MSS (maximum segment size). Prevents PacketShaper from changing the MSS on large WAN links.	384000 bps	0	512000
trafficIsAsymmetric By turning on this setting, PacketShaper will automatically assume all flows are asymmetric and stop TCP Rate Control. In topologies where there are a large percentage of asymmetric flows, this may be more efficient than attempting to apply regular rate control. In addition to disabling rate control, turning on this setting disables all layer 7 classification activities (PacketShaper must see traffic in both directions in order to classify layer 7).	0 (off)	0 (off)	1 (on)
uiDefaultSky The user interface that appears after logging in to the browser interface: the Sky UI or the original (legacy) user interface.	1 (Sky)	0 (legacy)	1 (Sky)
userEventExtSnmpVersion Enable/disable the extended SNMP trap for user events. When this variable is turned on, there will be an additional field in the trap that indicates the type of situation that triggered the trap. The field indicates *violated* (when the threshold was exceeded) or *rearm* (when the re-arm value was crossed).	0 (off)	0 (off)	1 (on)
userEventMaxDefinitions The maximum number of events that can be user-defined	32	32	128
userEventMaxRegistrations The maximum number of events that can be registered	32	32	128
userSessionIdleTimeout This variable controls how many seconds it takes for an unauthenticated login session to get purged from the system. You might need to increase this value if the session times out before PacketShaper can authenticate a login password, for example, when there is latency on the network or they are using a RADIUS or TACACS implementation. Note that this variable does not apply to idle sessions that have already been authenticated–just new sessions that have not yet been authenticated.	30	30	360 (6 min)

Command Change History

Release	Modification
11.7.1	**enableJumboFrame** variable introduced
11.6.2	**padVlanPkt** variable introduced
11.1.1	Removed variables for unsupported features
9.2.2	**userSessionIdleTimeout** variable introduced
8.6.1	**discoveryThresholdUrlCategories** variable introduced
8.5.3	Default setting for **PolicyFlowLimitForAllClasses** was changed to off
8.5.1	**uiDefaultSky** variable introduced
8.3.1	**diffservClassSortPref** variables introduced **browserHttpAcceleration** variable removed
8.2.0	**autoCreateSameSide**, **tcpClipInitialWindow** variables introduced
8.0.1	**userEventExtSnmpVersion** variable introduced
8.0.0	The following variables were introduced: **flowRecordsSendIntermediate**, **flowRecordsIntermediateTimeout**, **flowRecordsResetCounter**, **enableLatency**, **enableWinnyClassification**

setup web-proxy

Enable/disable the web proxy feature. When this feature is enabled, the configured web-proxy server handles the WebPulse requests, category map downloads, heartbeat emissions, support status updates, and image updates. Because some PacketShaper features (such as URL categories) access external servers on the Internet, you must configure an explicit web proxy for the PacketShaper if your company's security policy requires that all outbound traffic go through a proxy. The topology should look like this:

LAN Switch-->PacketShaper-->Proxy--->Router-->Internet

```
setup web-proxy on|off
```

Command Change History

Release	Modification
8.6.1	**setup web-proxy** command introduced

setup web-proxy default

Returns web proxy settings to their defaults so that they can be inherited from the parent configuration.

```
setup web-proxy default
```

Command Change History

Release	Modification
8.6.3	**default** option introduced
8.6.1	**setup web-proxy** command introduced

setup web-proxy server

Configure the IP address and port of the web proxy server. All PacketShaper features that access external servers on the Internet will go through the proxy server. This server handles WebPulse requests, category map downloads, heartbeat emissions, support status updates, and image updates.

```
setup web-proxy server <ip-address>:<port>
```

The *<ip-address>* of the web proxy server must be reachable by the PacketShaper.

Example:

```
setup web-proxy server 10.9.66.12:8000
```

> **Note:** You may need to add this server's IP address to the firewall's white list, if required.

Command Change History

Release	Modification
8.6.1	**setup web-proxy** command introduced

setup web-proxy show

Display the web proxy settings.

```
setup web-proxy show
```

Sample output:

```
Web Proxy Settings
        Status : on
     Server IP : 10.9.66.12
   Server port : 8000
```

Command Change History

Release	Modification
8.6.1	**setup web-proxy** command introduced

setup webpulse

Enables/disables the WebPulse enhanced classification feature. When WebPulse is enabled, the PacketShaper can classify by URL categories, web applications, and web operations. It is enabled by default.

```
setup webpulse on|off
```

When you enable WebPulse, you will have the ability to:

- create classes based on specific URL categories, and PacketShaper will then classify web traffic that corresponds to each of these categories into the appropriate class
- auto-discover category classes (you must first create an "all categories" class and enable class discovery)
- create classes based on specific web applications and operations
- auto-discover web application classes

When WebPulse is turned on, PacketShaper will send URL queries to the WebPulse cloud service; the service will look up the URL in its extensive database to find the categories, web applications, and/or actions associated with the URL, and send the response back to PacketShaper for classification. If WebPulse is disabled, the PacketShaper will look up the URL in its cache; if the URL is not in the cache, PacketShaper won't be able to classify the web traffic into a category.

The PacketShaper S500 has the option of using WebPulse Express, a feature that downloads the URL database to the appliance. When WebPulse is enabled and Express is configured, URL queries don't need to be sent to the WebPulse cloud service: they are looked up on the appliance, so classification is much faster. See setup webpulse express.

Command Change History

Release	Modification
9.2.1	**setup webpulse** command is introduced

setup webpulse cache-test

Looks up the specified URL in the PacketShaper's local cache and returns the category, web application, and operation names for the URL. Use this command to find out if a URL is in the cache or to see how a URL is being categorized.

```
setup webpulse cache-test <url>
```

where *<url>* is the URL to be looked up. Follow these guidelines when specifying the URL:

- The "www" is optional unless the server can host more than just Web, or if there is another subdomain you want to look up.
- Enter the URL exactly as it appears in your browser; you can copy the URL from the browser and paste it into the command window.
- Include https:// to do a lookup for a URL hosted on port 443.

```
# setup webpulse cache-test facebook.com

Category: Social Networking
Web application: Facebook
Operation: Unknown

# setup webpulse cache-test groupon.com

No categorization for groupon.com found.
```

> **Note:** If the command responds with *No categorization for <url> found*, the specified *<url>* is not in the cache. You can use the setup webpulse cache-update command to query WebPulse for the category.

Command Change History

Release	Modification
9.2.1	**setup webpulse** command is introduced

setup webpulse cache-update

Sends a query to the WebPulse cloud service to look up the category, web application, and operation of the URL. The URL and its category ID, web application ID, and operation ID are then placed into the URL cache, overwriting any existing entry. If the setup webpulse cache-test command failed because the URL was not in the cache, you can use the **cache-update** command to find the category, web application, or operation. This command is also useful if you suspect a change in categorization, such as a recent malware notification.

```
setup webpulse cache-update <url>
```

where *<url>* is the URL to be looked up. Follow these guidelines when specifying the URL:

- The "www" is optional unless the server can host more than just Web, or if there is another subdomain you want to look up.
- Enter the URL exactly as it appears in your browser; you can copy the URL from the browser and paste it into the command window.
- Include https:// to do a lookup for a URL hosted on port 443.

If the command responds with *<<Unrated>>*, the specified *<url>* is not in the WebPulse database. Check your spelling and try again.

Examples:

```
# setup webpulse cache-update craigslist.org

Category: Shopping
Web application: Craigslist
Operation: Unknown
```

```
# setup webpulse cache-update youtube.com

Category: Open/Mixed Content, Audio/Video Clips
Web application: YouTube
Operation: Unknown
```

Command Change History

Release	Modification
9.2.1	**setup webpulse** command is introduced

setup webpulse discovery

Enables/disables discovery of specific URL categories, or all categories. By default, discovery is enabled for each category. You may want to disable discovery of the categories that you aren't interested in monitoring or controlling. By doing so, you will not clutter your traffic tree with classes you don't need to track.

```
setup webpulse discovery on|off|inherit <category_name>|all
```

where *<category_name>* is the exact name of the category. To see a list of category names, use the **setup webpulse show categories** command. If the category name has a space, enclose the name in quotes.

If the unit is in shared mode, you can use the **inherit** option to inherit the category discovery settings from the parent configuration.

Example

To disable discovery of the Society/Daily Living category:

```
setup webpulse discovery off "Society/Daily Living"
```

To reenable discovery of all categories:

```
setup webpulse discovery on all
```

Command Change History

Release	Modification
9.2.1	**setup webpulse** command is introduced

setup webpulse express

WebPulse Express initiates the download of the entire WebPulse database to a PS-S500 appliance, and regularly downloads incremental updates to the database to keep the local content current. By having a local copy of the database on the appliance, the process of URL categorization and web application identification is significantly faster.

```
setup webpulse express set {username | password | interval} <value>
```

Contact your Symantec sales representative for the **username** and **password** to use with WebPulse Express.

The **interval** defines how frequently the PacketShaper checks for updates to the WebPulse database. The interval value is a positive integer, defined in seconds; the default value is 60 seconds.

After WebPulse verifies your credentials, the download of the 500MB database will begin. Use the **setup webpulse express show** command to monitor the status of the download. During the download process, the output will say *Loading the Database*. When download is complete, it will say *Database is loaded*, indicating the date and time of the download. If there was a problem loading the database (for example, the user name and password weren't configured properly), the output will say *Database is not loaded*.

The WebPulse Express username, password, and interval must be specified in separate commands, as shown in the examples below.

Examples

```
setup webpulse express set username bcwfname
setup webpulse express set password *****
setup webpulse express set interval 120
```

Additional Information:

- The WebPulse Express feature is available on PS-S500 appliances with 64 GB of RAM. Some models may have sufficient memory already; others may require the purchase and installation of a memory upgrade kit. To see if your appliance has sufficient memory, issue the " version " on page 1400command.
- WebPulse Express is not available on PS-S200 and PS-S400 models.

- Symantec will email you a username and password for WebPulse Express. Until WebPulse Express is configured properly with the correct user name and password, WebPulse classification will still function but will use the database in the cloud.
- Use the setup webpulse on command to enable URL categorization and web application classification.

To see the status of WebPulse Express, use the following command:

setup webpulse express show

```
Database is loaded at 2015-01-16 14:33:02.
Current database version is 350160400.
Database will update every 60 seconds.
Username: BCWF-DEC2814
Password: ******
```

Command Change History

Release	Modification
11.3.1	WebPulse Express feature is introduced

setup webpulse map-download

Initiates a download of the latest WebPulse mapping files for identifying URL category names, web application names, and web operation names.

```
setup webpulse map-download
```

The WebPulse maps support the following features:

- **URL categorization:** The category map associates URL category names with numeric IDs. PacketShaper refers to the map to look up the category name after WebPulse has assigned a category ID to a flow.
- **Web application identification:** The web application map associates web application names with numeric IDs, and lists the operation IDs applicable to the application. PacketShaper refers to the map to look up the application name after WebPulse has assigned an application ID to a flow. In addition, PacketShaper uses the map to determine which operations are applicable to an application, for example, to create a class for a specific operation for a web application (such as Facebook-Post_Messages, Gmail-Upload_Attachment).

Additional Information

- When you enable WebPulse, the latest map files are automatically downloaded to your PacketShaper. Thereafter, PacketShaper automatically downloads the maps every day and after a device reset, regardless of whether the files have changed or not.
- You will see a *Success* message after the download completes. The map files are downloaded, replacing the current files (even if the file hasn't changed). The version number is also indicated.
- The map version is incremented when a category, web application, or operation has been added, renamed, or deleted.

Command Change History

Release	Modification
9.2.1	**setup webpulse** command is introduced

setup webpulse reset

Clears the URL cache on the PacketShaper, deletes the cache backup, and returns all WebPulse settings to their defaults. If WebPulse Express is in use, the reset command removes the on-box database and clears the Express username and password. This command returns the feature to its factory default settings: WebPulse is enabled and discovery is enabled for all URL categories.

```
setup webpulse reset
```

> **Note:** Because this command deletes the cache as well as its backup copy in the PacketShaper 9.1026/urlcat/cache folder, use this command with caution. You will have an opportunity to confirm the reset after you enter the command.

Command Change History

Release	Modification
9.2.1	**setup webpulse** command is introduced

setup webpulse show

Displays additional information about WebPulse statistics, caches, service points, and so forth.

```
setup webpulse show categories|statistics|cache|service|applications|operations
```

Option	Description
categories	Display URL category names, IDs, discovery state, and hits
statistics	Display WebPulse statistics: how many flows weren't processed due to load or because the flow ended, number and speed of queries to the WebPulse database, number and speed of queries that used the Dynamic Real-Time Rating (DRTR) service.

Option	Description
cache	The URL caches contain URLs and their category, application, and operation IDs. This command displays statistics about the URL caches: number of hits and entries in the domain, directory, and filename caches. Also indicates when the cache was last backed up to the 9.1026/urlcat/cache folder. Examples of URLs that go in each type of cache: Domain cache: bluecoat.com Directory cache: nps.gov/yose Filename cache: cnn.com/forum/viewforum.php?f=1
service	Display WebPulse service point information. For each of eight service points, the output lists the IP address, number of hits and speed of the WebPulse Rating Service (RS) database and Dynamic Real-Time Rating (DRTR) requests. Also indicates the health of the service points. **Tip:** The fastest servers appear at the top of the list.
applications	Display name and supported operations for each web application
operations	List all supported web operation names

Examples:

```
setup webpulse show statistics


Unprocessed Flows

 Current Daily Unprocessed Due to Load: 0
    Average Daily Unprocessed Due to Load: 0
    Current Daily Unprocessed Due to Flow Ended: 89
    Average Daily Unprocessed Due to Flow Ended: 102


Service Points

    Current Daily Queries        : 161
    Current Daily Query RTT      : 48 ms
    Average Daily Queries        : 598
    Average Daily Query RTT      : 51 ms
    Current Daily DRTR Queries   : 3
    Current Daily DRTR Query RTT : 966 ms
    Average Daily DRTR Queries   : 0
    Average Daily DRTR Query RTT : 0 ms


setup webpulse show cache


Cache
    Daily Queries           : 1100
    Average Daily Queries   : 2417
    Cache Hits              : 926
    Average Daily Cache Hits: 1736
    Average Cache Efficiency: 71.83%
Domain Cache
    Cache Hits              : 872
```

```
    Average Daily Cache Hits: 1575
    Number of Entries       : 575
    Average Hourly Size     : 516
Directory Cache
    Cache Hits              : 54
    Average Daily Cache Hits: 161
    Number of Entries       : 930
    Average Hourly Size     : 863
Filename Cache
    Cache Hits              : 0
    Average Daily Cache Hits: 0
    Number of Entries       : 0
    Average Hourly Size     : 0
URL Cache Backup Status:
    Cache backup successfully updated on Tue Aug  3 14:56:20 2010
```

> **Note:** To show details about WebPulse Express, use the setup webpulse express show command.

Command Change History

Release	Modification
9.2.1	**setup webpulse** command is introduced

shutdown

Shut down the operating system and file system and power off the PacketShaper. This command is used to prepare the PacketShaper for transport.

```
shutdown
```

Command Change History

Release	Modification
11.1.1	Command introduced

synthetic add

Create a new synthetic transaction. Using synthetic transactions allows PacketShaper to initiate ICMP, web, or other TCP transactions at periodic intervals to verify the availability of critical hosts.

```
synthetic add <interval>[,<repeat>] [<id>] <url>
```

<interval>	Number of minutes between issuance of the transaction (the maximum interval is 1440)
[, *<repeat>*]	Number of times to issue the request on the established TCP connection (default is 1; the maximum is 100)
<id>	String that identifies the synthetic transaction; if omitted, PacketShaper will automatically create a unique eight-character ASCII ID for each transaction.
<url>	Type of transaction to issue, in the following format: *<type>*://*<host>*[:*<port>*][/*<path>*] where: *<type>* is **http**, **https**, **icmp**, **pop3**, **smtp**, **ftp**, **echo**, or **custom** **Note:** The *<type>* must be entered in lowercase. *<host>* is the IP address or DNS name of the host *<port>* is the TCP port number to connect to; the default varies by type (for example, for http the default is port 80) *<path>* is additional information necessary for the request (such as a folder or file name)

Additional information about each type:

- The **http** type will issue a GET request for the file specified by the *<path>* parameter. (The default port is 80.)
- The **https** type does an SSL handshake and issues a GET request for the file specified by the *<path>* parameter. (The default port is 443.)
- The **icmp** type sends a ping request to the designated server, using the ICMP-ECHO Protocol (RFC 792). Example: `synthetic add 17,2 "icmp://www.bluecoat.com"`
- The **smtp** and **pop3** types also do not send or receive mail; they issue a single command over the channel to elicit a response. (The default port is 25.)
- The **ftp** type will issue a single retrieve command (RETR) for the file specified in the *<path>* parameter but doesn t do any user authentication. (The default port is 21.)
- The **echo** type sends a string to the designated host and the host echos it back. TCP echo requires that the target host have an echo server process running and listening on port 7. The optional *<path>* argument has the format *<length>* [/*<fill>*] where *<length>* is the number of bytes to send on each request (the default is 512) and *<fill>* is a string to fill the request buffer. The *<fill>* string can be up to 511 bytes.

For example:

```
echo://test.domain.com/10/xyz
```

The above example sends requests containing xyzxyzxyzx (10 bytes).

- The **custom** type allows you to specify a series of requests to be sent alternatively for as many messages as requested by the *<repeat>* parameter. The request strings are separated by the | character. For example:

```
custom://my.test.com:25/HELO|MAIL FROM:<bob>|RCPT TO:<brett>|DATA|hey|.
```

The above example sends a simple message to a mail server on port 25 (the default port for SMTP).

Command Change History

Release	Modification
8.0.1	**icmp** *<type>* introduced
8.0.0	no change

synthetic delete

Delete a synthetic transaction.

```
synthetic delete <id>
```

where *<id>* is the identifying name specified in the synthetic add command. To view IDs of all synthetic transactions, use the synthetic show command.

When a synthetic transaction is deleted, its corresponding *Inbound/SyntheticTransactions* or *Outbound/SyntheticTransactions* traffic class is not deleted, so that measurement data can still be retrieved from that traffic class. Even after the synthetic transaction is deleted, network traffic may still be classified in that traffic class until the class is also manually removed.

synthetic options

Create traffic classes for the hosts specified in synthetic transactions. The classes will be created in the *SyntheticTransactions* class.

```
synthetic options create-classes show|on|off|default
```

The default value is **on**. If you have already created traffic classes for your critical hosts and you want synthetic transaction measurement data to be recorded in these classes, set this option to **off**.

> **Note:** If you use the CLI command **synthetic options create-classes on** to create traffic classes for synthetic transactions hosts and then later issue the **command synthetic options create-classes off** to turn off this option, any traffic classes already created for previous synthetic transactions will remain a part of the configuration's traffic tree.

When the **`synthetic option create-classes default`** command is issued for an individual PacketShaper or a PolicyCenter parent configuration at the top of the PolicyCenter configuration tree, this command restores the default **on** value. When this command is issued for a PolicyCenter child configuration, the child configuration clears its local setting and inherits the **synthetic option create-classes on|off** value from its parent configuration.

Command Change History

Release	Modification
8.0.1	Synthetic transaction classes created as children of Localhost class instead of in a SyntheticTransactions folder <*host*> cannot be an IP address
8.0.0	no change

synthetic show

Display information about synthetic transactions.

```
synthetic show
```

The output displays all the active synthetic transactions, when they are next scheduled to run, and a count of how many TCP connections have been attempted and were accepted.

Transaction ID	URL	Interval	Repeat	Next Scheduled	Attempts	Connections
shop1	http://www.cdnow.com	5	1	26-Jul-2001 14:30:38	0	0
st4	echo://10.10.10.10/monsters	1	5	26-Jul-2001 14:28:34	0	0
st2	http://www.lucent.com/minds/innovating/index.html	3	3	26-Jul-2001 14:30:00	6	5
st3	custom://my.test.com:80/HEY YOU	4	1	26-Jul-2001 14:31:17	5	4
st1	http://www.amazon.com	2	2	26-Jul-2001 14:30:04	9	9

sys

These diagnostic commands are intended to be used only under the guidance of Customer Support and are not covered in this guide.

sys limits

List the PacketShaper's configuration limits. For each object (such as classes, partitions, and policies), the **sys limits** output lists the maximum number of objects allowed, currently used, and remaining. For example, you can use this command to determine how many more classes you can create on your unit.

```
sys limits

Statically allocated objects          Current  Remaining      Total
 ---------------------------------------------------------------------
 Traffic classes                          800        224       1024
 Partitions                                 2        510        512
 Dynamic Partitions                         0      10000      10000
 Policies                                   4       1020       1024
 Matching rules                          1896       3224       5120
 Classes with top talkers/listeners         0         12         12
 TCP flows                                 36     199964     200000
 Other IP flows                            82      99918     100000
 Legacy flows                               8       9992      10000
 Concurrent Hosts                        6250      93750     100000
 MAC Cache Entries                        321      14679      15000
 Fragment Cache Entries                     0       8000       8000
 Command Contexts                           6         24         30

Dynamically allocated objects       Current  Potential     Total

---------------------------------------------------------------

Matching rule host references           7      10253      10260
 Host list DS entries                   0      24044      24044
 DNS names                             10      59064      59074
```

Note: "Potential" for each object is an estimate allocating all remaining dynamic memory to that object type.

The table below describes the items of interest in the **sys limits** output.

Object	Description
Statically Allocated Objects	
Traffic class	A logical grouping of traffic flows that share the same characteristics a specific application, protocol, address, or set of addresses.
Partition	A bandwidth pipe assigned to a given traffic class to protect or restrict the total bandwidth available to that class.

Object	Description
Dynamic partitions	A type of partition that automatically creates subpartitions on the fly as users become active in a traffic class. This capability allows service providers or enterprise customers to guarantee each user a minimum amount of bandwidth at all times
Policy	A rule assigned to a given traffic class that defines how a single flow will be handled during bandwidth allocation.
Matching rule	A set of characteristics that identifies a specific traffic type.
Classes with top talkers/listeners	An identified traffic class for which PacketShaper has been configured to record the host names or IP addresses of the devices transmitting the greatest amounts of traffic (the "talkers") or receiving the greatest amounts of traffic (the "listeners").
TCP flows	Unique sessions using Transmission Control Protocol.
Other IP flows	Unique sessions using non-TCP Internet protocols.
Legacy flows	Traffic using a non-TCP/IP protocol, often encapsulated in a TCP or IP wrapper.
Concurrent hosts	The Current count for concurrent hosts is an indication of how many host addresses have been learned by the system, that is, the number of entries currently in the host database. Entries in the host database are not periodically aged out or cleared from memory – instead, they are reused when needed for new hosts. Therefore a value of 0 in the Remaining column does not mean your unit can't accommodate any more hosts – new hosts will simply replace hosts that are no longer active. It is normal and expected to see the number of concurrent hosts at its limit.
Dynamically Allocated Objects (maximums can vary, depending on the amount of remaining dynamic memory)	
Matching rule host references	Unique domain names in matching rules.
Host list DS entries	A set of IP addresses and/or DNS names that traffic class matching rules can reference.
DNS names	Unique domain names used in PacketShaper configuration (in matching rules, configured SNTP time servers, configured RADIUS authentication and accounting servers, etc.) Use the dns names CLI command to see a list of domain names in use.
Classes with worst clients/servers	Clients or servers that have the highest percentage of transactions exceeding the total delay threshold.

tenant add customer

Assign a numeric ID to a customer. Service providers can assign an ID to each customer. Even if your organization doesn't have multiple customers, you must create an ID so that you can then define the default Acceptable Use Policy for the company.

```
tenant add customer id <cust_ID>
```

The *<cust_ID>* must be a positive integer, eight digits or less. Each customer must have a unique ID.

Example

```
tenant add customer id 1
Customer added.
```

Command Change History

Release	Modification
11.8.1	Command introduced

tenant add customer aup

Define the default Acceptable Use Policy (AUP) for a customer. The AUP message will appear in the user's browser window when he/she attempts to access a website with an HTTP URL. (If an HTTPS URL is entered, the connection will be dropped.)

`tenant add customer <cust_ID> aup "aup/<filename>.htm"`

To specify an AUP, reference the name of an HTML file that contains the policy message. PacketShaper comes with a sample file that you can either use as is or customize with your own message and customer logo. The sample file, named aup_example.htm is located in the 9.258/customer/aup directory. To customize the file, use the File Browser to view and save the sample file to a local workstation; after customizing the HTML, you can upload the new file to the same directory. If you referenced an image in the HTML file, you also need to upload the image file.

Example

To define the AUP for customer 1:

```
tenant add customer 1 aup "aup/testcorp_1.htm"
Customer AUP setting updated.
```

Notes

- This AUP file is served to all hosts in the customer's managed subnet (as defined by the tenant set start-ip CLI command), unless a different AUP file is defined for the site the host is in. The site AUP file overrides the customer file.
- Use the tenant show customer command to view the default AUP filename for a customer.
- The *<filename>* must have an .htm file extension (.html will not work).
- The AUP file contents and filename are not validated at the time you configure the customer AUP.

Command Change History

Release	Modification
11.8.1	Command introduced

tenant add site

Assign a numeric ID to a site. If a customer has multiple subnets (each with a different Acceptable Use Policy), you will create a unique site for each subnet.

`tenant add site id <site_ID> <cust_ID>`

The *<site_ID>* is a positive integer, eight digits or less. Each site must have a unique ID and the value cannot be a duplicate of a site in another customer. In other words, you cannot have a site 1 in customer 1 and in customer 3.

Example

To create a site ID 1 for customer 2:

```
tenant add site id 1 2
Site added.
```

After creating the site ID, you can assign attributes to the site, such as specifying the subnet and custom AUP. See tenant add site <attribute>.

Command Change History

Release	Modification
11.8.1	Command introduced

tenant add site *<attribute>*

Set the attributes for a site. At a minimum, you need to define the subnet and captive portal IP for the site. Optionally, you can define an Acceptable Use Policy (AUP) that overrides the default AUP for the customer. You can also assign a descriptive name to the site (such as *First Floor Conference Rooms*) and create a whitelist of IP addresses that are exempt from having to sign in to the captive portal.

`tenant add site <site_ID> <attribute>`

where *<attribute>* is one of the following:

`aup "aup/<filename>.htm"`	Define the site AUP HTML filename; this file overrides the customer AUP defined with the tenant add customer aup command. PacketShaper comes with a sample file that you can either use as is or customize with your own message and customer logo. The sample file, named aup_example.htm is located in the 9.258/customer/aup directory. To customize the file, use the File Browser to view and save the sample file to a local workstation; after customizing the HTML, you can upload the new file to the same directory. If you referenced an image in the HTML file, you also need to upload the image file.
`dst-exemption <IP or subnet>`	Specify a destination IP address or subnet that is exempt from the captive portal; if a user goes to a destination IP address that is exempt, the AUP message will not display. Each destination IP address or subnet must be specified in a separate command. Example subnet: *192.168.23.0/24*
`name <site_name>`	Assign a descriptive name (up to 127 characters) to the site. Enclose the name in quote marks.
`redirect <IP>`	Define which captive portal IP address to use for the site. This IP address must in the list of captive portal IPs defined with the tenant set captive add command.
`src-exemption <IP or subnet>`	Specify an IP address or subnet that is exempt from the captive portal; the AUP message will not display for exempt source IP addresses. Each source IP address or subnet must be specified in a separate command. Example subnet: *192.168.23.0/24*
`subnet <IP> <mask>`	Define the subnet for the site. This subnet must be within the scope defined by the tenant set start-ip and tenant set netmask commands and must be non-overlapping with other sites.

Examples

```
tenant add site 1 subnet 10.100.1.0 255.255.255.0
WiFi subnet added.
tenant add site 1 redirect 1.2.3.4
Redirection IP updated.
tenant add site 1 name "Conference Rooms"
Site name setting updated.
tenant add site 1 aup "aup/site1_aup.htm"
Site AUP setting updated.
tenant add site 1 src-exempt 10.100.1.24
```

```
Source IP address exemption added.
```

Notes

- Use the tenant show site command to view the attributes assigned to a site.
- Use the tenant remove site command to remove attributes assigned to a site.

Command Change History

Release	Modification
11.8.1	Command introduced

tenant enable

Enable/disable the captive portal feature. Captive portal is disabled by default, and must be disabled while configuring the basic settings of the feature (those defined with tenant set).

```
tenant enable | disable
```

Command Change History

Release	Modification
11.8.1	Command introduced

tenant remove

Remove a captive portal customer, site, or site attribute.

`tenant remove customer id <cust_id> | site id <site_id> | site <site_ID> <attribute>`

> **Caution:** Removing a customer deletes all sites associated with the customer and you will not have the opportunity to confirm or undo the deletion. Use the tenant show customer command first to confirm that you are removing the intended customer.

The following attributes can be removed from a site: `subnet`, `src-exemption`, `dst-exemption`, `aup`. You cannot remove the `redirect` attribute; you must assign one of the captive portal IP addresses that were defined with the tenant set captive add command.

Examples

```
tenant remove site 1 src-exemption
Source IP address exemption removed.
tenant remove site id 1
Site removed.
```

Command Change History

Release	Modification
11.8.1	Command introduced

tenant set

Configure the captive portal feature.

`tenant set start-ip <IP> | netmask <mask> | captive {add <IP> | del <IP>}`

`start-ip <IP>`	Set the starting IPv4 address of the subnet to be managed. The guest WiFi network must be on a non--overlapping IPv4 subnet from the corporate network.
`netmask <mask>`	Enter the mask of the subnet to be managed (for example, 255.255.0.0).
`captive add <IP>`	Specify the IPv4 address that represents an on-box PacketShaper web server. This IP address is used internally to perform web redirection (for example, to redirect from the user's requested URL to the Acceptable Use Policy page). The IP can be any valid address that is not in use by other hosts on the network. You can add up to eight captive portal IP addresses, with each one serving a different AUP (for example, customized to different customers).
`captive delete <IP>`	Remove a captive portal IP address.

Examples

```
tenant set start-ip 10.100.0.0
tenant set netmask 255.255.0.0
tenant set captive add 1.2.3.4
```

Notes

- Use tenant show status to verify captive portal settings.
- Use tenant enable to turn on the captive portal feature, after all settings (including customers and sites) have been defined.
- You cannot configure the captive portal settings while the feature is enabled. Before modifying settings, use the **tenant disable** command to turn off the feature.

Command Change History

Release	Modification
11.8.1	Command introduced

tenant show

View settings related to the captive portal feature.

`tenant show status | customer <cust_ID> | site <site_ID>`

`status`	Display the captive portal source IP subnet and IP address(es) of configured captive portal web servers on PacketShaper.
`customer <cust_ID>`	View a customer's settings where *<cust_ID>* is the value assigned to the customer.
`site <site_ID>`	View a site's settings where *<site_ID>* is the value assigned to the site.

Examples

```
tenant show status
Captive Portal is OFF
Source IP subnet is 10.167.9.0/24
Captive Portal redirection IP:
10.9.43.220
tenant show customer 2
-------------------------------------------------
Customer ID: 2
Customer Sites:
site1 (id: 1)
Customer Guest WiFi AUP: aup/customer2_aup.htm
-------------------------------------------------
tenant show site 1
-------------------------------------------------
Site ID: 1
Site Name: site1
Customer ID: 2
Site Guest WiFi Source exemption:
10.167.9.2 / 255.255.255.255
Site Guest WiFi Destination exemption:
Site Guest WiFi Subnet: 1.1.1.0 / 255.255.255.0
Site Guest WiFi AUP: aup/site1_aup.htm
Site redirection IP: 10.9.43.220
-------------------------------------------------
```

Command Change History

Release	Modification
11.8.1	Command introduced

traffic active

Display the current, maximum, and possible number of sessions for TCP, UDP, and Legacy traffic types. This command is a valuable tool for determining how close the unit is to reaching its capacity. It also gives a histogram of the number of host

entries in various time buckets (based on idle time).

```
traffic active
```

	TCP	UDP	Legacy	Total
Flows (Current):	13	32	9	54
Flows (Maximum):	27	59	13	67
Flows (Possible):	50000	25000	5000	80000

Host Entries Histogram (based on idle time):

<1s	<1min	<2min	5min	<10min	>10min
5	13	26	5	7	10184

Type of Flow	Description
TCP flows	Unique sessions using Transmission Control Protocol.
UDP flows	Unique sessions using non-TCP Internet protocols, such as UDP.
Legacy flows	Traffic using a non-TCP/IP protocol, often encapsulated in a TCP or IP wrapper.

The value listed for maximum flow is the maximum number that has been displayed when the **traffic active** command has been executed (since the last reset). In other words, maximum numbers are only recorded when this command is run, so the maximum counts are representative rather than authoritative values.

The possible flows represent the unit's maximum number of concurrent flows allowed on the unit. PacketShaper can support more flows than the indicated number, but these figures represent the ideal maximums for producing optimal results.

traffic bandwidth

Display bandwidth utilization for a partition.

```
traffic bandwidth [<tclass>] [clear]
```

[*<tclass>*]	The root traffic class of the partition to display. If you do not specify a traffic class, the outbound partition statistics are displayed.
[clear]	Clears the accounting data

When traffic shaping is turned off, the **traffic bandwidth** command displays the aggregate usage summary for both the inbound and outbound directions.

Example: The inbound/http class has a 500k partition, burstable to 1Mb. The class also has a 0k rate policy with priority 3. Here is the output of the **traffic bandwidth inbound/http** command:

inbound partition HTTP				
Programmed min bandwidth		500k	max bandwidth	1.0M
Adjusted min bandwidth		500k	max bandwidth	1.0M
Local		500k		
reserved rate	0	unreserved rate/limit	0/0	
reserved peak	0	unreserved peak	0	
ignored rate	2220			
ignored peak	12.0M			

current guaranteed rate	0	excess rate	0
OVERalloc'ed guaranteed	0	excess	0

Gain: 0.00 Compensation 0

Excess Rate

priority	0	1	2	3	4	5	6	7
demand	0	0	0	0	0	0	0	0
% satisfied	100	100	100	100	100	100	100	100

Priority Traffic

priority	0	1	2	3	4	5	6	7
pkts relayed	0	0	0	0	0	0	0	0
exceptions	0	0	0	0	0	0	0	0
anticipations	0	0	0	0	0	0	0	0

Refer to the following list for **traffic bandwidth** details:

min bandwidth	The partition size in bits per second (e.g., 500k bps); the Programmed (initially-configured) size and Adjusted (actual) size are displayed. Note: The Programmed and Adjusted values may differ from one another in a hierarchical partition. Because a child partition is a percentage of the parent partition, if a parent gets less bandwidth, the child will also get proportionally less.
max bandwidth	Maximum (burstable) partition size; the Programmed (initially-configured) size and Adjusted (actual) size are displayed.
reserved rate	The total bandwidth currently in use by rate-based traffic
reserved peak	The peak bandwidth usage by rate-based traffic
unreserved rate/limit	The total bandwidth currently in use by priority traffic
unreserved peak	The peak bandwidth usage by priority-based traffic
ignored rate	The current bandwidth in use by ignore-policy traffic classes and uncontrolled traffic
ignored peak	The peak bandwidth usage by traffic classes with the ignore policy
current guaranteed rate	The total current guaranteed rate usage in bits per second
excess rate	The total current excess rate usage in bits per second
OVERalloc'ed guaranteed	The amount of guaranteed rate that currently is over-allocated
Excess Rate priority	Priority levels (0-7)
Excess Rate demand	The excess rate demand at each priority level
Excess Rate % satisfied	The percentage of excess rate demand that is currently being satisfied at each priority level

Priority Traffic priority	Priority levels (0-7)
Priority Traffic pkts relayed	Number of packets counted at each priority level
Priority Traffic exceptions	PacketShaper uses a rate anticipation mechanism to shape the rate of non-TCP traffic. When this mechanism fails to keep the desired rate from being exceeded, a rate exception is counted.
Priority Traffic anticipations	When PacketShaper determines that the desired rate is in danger of being exceeded, an exception anticipation event is counted.

traffic flow

Display summary information about some or all currently active TCP connections and/or UDP sessions.

`traffic flow -tIo`	TCP overview of non-idle flows
`traffic flow -uIo`	UDP overview of non-idle flows
`traffic flow -h`	lists help with all options

Example:

```
traffic flow -t
Num TCP Flows total   =    115 (all classes)
InAddr                    OutAddr                  Idle Svc
--------------------------------------------------------------------------------
2001:db8:1234:5678::1     2001:db8:1234:5678::2      1s  FTP-Cmd-Clear-IPv6
2001:db8:1234:5678::1     2001:db8:1234:5678::2      2s  SSL-IPv6
2001:db8:1234:5678::1     2001:db8:1234:5678::2      0s  FTP-Cmd-Clear-IPv6
10.9.45.20                10.9.60.138                2s  HTTP
10.9.60.109               10.9.60.110                1s  SSL
10.9.60.109               10.9.60.110                2s  FTP-Data-Clear
2001:db8:1234:5678::1     2001:db8:1234:5678::2      1s  SSL-IPv6
```

To display more detailed information, use:

For TCP flows: `traffic flow -t[aAcCfhiILmnNoOpPsSUvVxX]`

For UDP flows: `traffic flow -u[aAcCfhiInoOpPvVxX]`

-a	all (same as -pifvs: could wrap around on screen)
-A <*addr*>	address (only show conns for specified address)
-c <*class*>	class (only show info for specified class name)
-C	class (show class names)
-f	flags (connection flags)

-g <*addr*>	owner (only show connections for standby owner 'owner')
-G	owner (show standby owner)
-h	help (show this help)
-i	idle (show idle time)
-I	non-idle (don't show flows idle for one minute or more)
-L	license (display license state of TCP flows)
-m	mss (show mss info - tcp)
-n <*num*>	num (show up to <*num*> flows)
-N	state (flows not in connected state - tcp)
-o	overview (show summary information only)
-O	overview (include summary information)
-p	ports (show port numbers)
-P <*port*>	port (show only port <*port*>)
-s	state (display state - tcp) When the -s flag is used, the output displays the following connection states in the **S** column: I=idle – The connection is idle. C=connecting – Client is trying to establish a TCP connection with server; client sends the first CONNECT/SYN segment to server. W=ackWait – After sending the CONNECT/SYN segment, the client is waiting from ACK from the server. X=dataXfer – Both sides of the TCP connection have exchanged SYN/ACK segments. H=halfDiscon – The client/server has received a first FIN segment from the application and bandwidth resources were released for that half of the connection. The other side of the TCP connection has not received the final FIN segment. D=disconnected – The final FIN is received. F=fading – An event (such as a Ctrl+C from the application) causes the TCP connection to abort.
-S <stat>	state (only display conns in state <*stat*> - tcp)
-t	tcp (show TCP flows)
-T	tcp rate (show the current TCP rate)
-u	non-TCP IP flows, such as UDP flows, RSVP, ICMP
-U	unique host pairs for a given class; only valid with -c
-v	service (display service info)
-V <*serv*>	service (display only services matching <*serv*>)

-x
expanded (show full class names)

Note: Must be used with either the -c or -C option.

-X
expanded (show full class names in multi-line output format)

Note: Must be used with either the -c or -C option.

When the **-f** option is used, the Flags column in the output displays one or more of the following:

A	Asymmetric
B	bad seq or ack
C	Fully classified
D	data, no SYN or SYN-ACK
I	Inbound closed
O	Outbound closed
s	SYN, no SYN-ACK
S	SYN-ACK, no SYN
T	Tentative
W	Web

> **Note:** If the service has not been determined, noted by a dash (-) in the Svc column, the side of the addresses may also be undetermined. For more information on which side of the unit a particular host is on, use the **host show** command.

The **traffic flow -tL** command puts two columns in the output: LI (representing the inbound part of the TCP flow) and LO (outbound part of the flow). A + indicates flows have been granted a license, a - indicates flows have been denied a license.

Note that some flows do not completely shut down, and are therefore listed until the unit is reset. Therefore, the -t or -u option, combined with the l option, provides a list of non-idle TCP or UDP flows. For example:

traffic flow -tlpc inbound/http

```
Num TCP Flows total = 3 (class HTTP)
InAddr        Port OutAddr         Port Idle  ClasI  ClasO  Svc
-------------------------------------------------------------------------
10.10.254.249    1119 207.158.237.171     80              22m  HTTP   HTTP HTTP
 10.10.254.249    1120 207.158.237.171     80              22m  HTTP   HTTP   HTTP
 10.10.254.249    1105 207.158.237.171     80              44s  HTTP   HTTP   HTTP
```

```
traffic flow -tCx
Num TCP Flows total   =      6 (all classes)
InAddr           OutAddr         Idle ClasI ClasO Svc
-------------------------------------------------------------------------------
172.21.19.102   10.1.1.46        8m  /Inbound/Default /Outbound/Default Telnet-Clear
```

```
172.21.1.39     10.100.99.30     10s /Inbound/Default /Outbound/Default KaZaA-Cmd
172.21.1.39     10.1.1.45        5m  /Inbound/NetBIOS-IP /Outbound/Default NetBIOS-IP-SSN
172.21.1.39     10.1.1.20        15s /Inbound/Microsoft-ds /Outbound/Default Microsoft-ds
172.21.1.39     10.1.1.18        8s  /Inbound/Default /Outbound/Default -
172.21.1.39     10.100.99.30     8s  /Inbound/Default /Outbound/Default KaZaA-Cmd


traffic flow -tCX
Num TCP Flows total   =      6 (all classes)
InAddr           OutAddr         Idle Svc
-------------------------------------------------------------------------------
172.21.19.102   10.1.1.46        8m  Telnet-Clear
Inbound Class: /Inbound/Default
Outbound Class: /Outbound/Default
172.21.1.39     10.100.99.30     21s KaZaA-Cmd
Inbound Class: /Inbound/Default
Outbound Class: /Outbound/Default
172.21.1.39     10.1.1.45        5m  NetBIOS-IP-SSN
Inbound Class: /Inbound/NetBIOS-IP
Outbound Class: /Outbound/Default
172.21.1.39     10.1.1.20        8s  Microsoft-ds
Inbound Class: /Inbound/Microsoft-ds
Outbound Class: /Outbound/Default
```

To view a list of unique host pairs for a traffic class (Inbound/Default):

```
traffic flow -taUc inbound/default
Num unique host pairs total   =     1 (class Default)
InAddr            OutAddr         # of flows
-------------------------------------------
10.9.50.157       10.9.50.75        10
```

To see if a host is being classified correctly in the expected class:

```
tr fl -tupXICA 209.210.203.33

Num TCP Flows total   =      1 (all classes)

InAddr           Port  OutAddr          Port Idle Svc
-------------------------------------------------------------------------------
192.168.0.7      4721 209.210.203.33     80  19s HTTP
Inbound Class: /Inbound/HTTP
Outbound Class: /Outbound/HTTP

Num UDP Flows total   =      0 (all classes)
```

Command Change History

Release	Modification
11.6.1	**-g** and **-T** parameters added
8.6.1	**-U** parameter added

traffic guaranteed

Display guaranteed rate utilization for a traffic class subtree.

```
traffic guaranteed [<tclass>] [clear]
```

[*<tclass>*]	The root traffic class of the subtree display. The class' explicit path is required *only* if the class name itself is not unique. If you do not specify a traffic class, the guaranteed rate information for the entire tree is displayed.
[clear]	Clears the accounting data

The command output displays a list of all child classes and the following associated information for classes with guaranteed rate policies:

- Current number of users
- Peak number of users
- Guaranteed bandwidth in bps
- Number of guaranteed rate failures

traffic history

Display recent traffic flows for a specific host or traffic class.

```
traffic history recent|find <name>
```

recent	Lists recent flows for a specified traffic class. The output includes the date, time, IP address, port number, and URL of each flow in the specified class.
find	Lists recent flows for a specified host. The output lists each class that the specified host uses, as well as the date, time, service name, IP address, port number, and URL of each flow in the class.
<name>	With the **recent** argument, *<name>* is the traffic class name. With the **find** argument, *<name>* is the IP address or name of the host to be tracked.

Examples

The **traffic history find** command is useful for determining the servers that a specified client IP address is transferring data with, or the clients that are retrieving data from a specific server. It can also be used to determine exactly what type of network applications a specified PC is using.

traffic history find 10.10.1.6

-----(/Outbound/rsh)-----

```
07-Jan-2005 10:53:25     rsh
 192.21.1.26     1023  raltman-t23.example.com
 10.10.1.6       514  test2.example.com

-----( /Inbound/rsh )-----

07-Jan-2005 10:53:25     rsh
 192.21.1.26      1023  raltman-t23.example.com
 10.10.1.6        514  test2.example.com
```

The **traffic history recent** command is useful for analyzing the type of traffic that is falling into a Default class, such as Inbound/Default in the following example.

traffic history recent inbound/default

```
-----( /Inbound/Default )-----

07-Jan-2005 13:01:19      UDP
 192.21.1.26       3288  example-40vp63
 10.100.10.30      2687  mail.example.com
 07-Jan-2005 12:59:53      UDP
 192.21.1.26       3299  example-40vp63
 192.21.0.20        389  dc-dev.example.com
 07-Jan-2005 12:56:14      UDP
 192.21.255.255    7741
 192.21.31.251     32808  opslab.example.com
 07-Jan-2005 12:42:16      TCP
 192.21.0.25       9100
 10.10.100.24      1995  phogle.example.com
 07-Jan-2005 12:33:19      UDP
 192.21.1.26       2967  example-40vp63
 10.10.10.18       2967  test.example.com
 07-Jan-2005 11:01:29      UDP
 192.21.1.26      38293  example-40vp63
 10.10.10.89       1046  test.example.com
 07-Jan-2005 10:51:54      HTTP
 192.21.1.26       2606  example-40vp63
 216.148.237.145     80  a216-148-237-145.deploy.akamaitechnologies.com
 07-Jan-2005 10:51:54      HTTP
 192.21.1.26       2607  example-40vp63
 216.239.53.104      80
 07-Jan-2005 10:51:54      HTTP
 192.21.1.26       2611  example-40vp63
 128.242.107.114     80  vrp1.sjc.xpc-mii.net
 07-Jan-2005 10:44:53      UDP
 255.255.255.255    631
 192.21.1.34        631
```

traffic licenses

Show current license usage for classes that have had the number of TCP flows limited with the class licenses command.

```
traffic licenses
```

Sample output:

```
    Traffic                        --- Licenses ---
     Class                         Total     In Use
-------------------------------------------------------
Inbound
 HTTP                                40         4
```

Another way to see the number of active TCP flows for a class is with the traffic flow command (for example, *traffic flow -tIpc inbound/http*). The output for this command will not only show the total number of licenses available and in use, but will also display details about each flow.

```
traffic flow -tIpc inbound/http
Num TCP Flows total   =      7 (class HTTP, licenses=7/40)
InAddr           Port  OutAddr          Port Idle ClasI ClasO Svc
-------------------------------------------------------------------------------
192.168.0.4      2694 64.236.43.54       80  39s HTTP HTTP HTTP
192.168.0.4      2676 216.148.237.36     80  48s HTTP HTTP HTTP
192.168.0.4      2671 64.12.174.57       80  40s HTTP HTTP HTTP
192.168.0.4      2687 64.236.43.54       80  39s HTTP HTTP HTTP
192.168.0.4      2689 64.236.43.54       80  39s HTTP HTTP HTTP
192.168.0.4      2670 64.12.174.57       80  39s HTTP HTTP HTTP
192.168.0.4      2690 64.236.43.54       80  39s HTTP HTTP HTTP
```

Note that the **traffic flow** command does not require that a limit be set with the **class licenses** command it will show the total number of active flows in any class you specify.

traffic statistics

Show statistics about current network traffic activity, dropped packets, and packet size distribution.

```
traffic statistics
```

Sample output:

```
                          Traffic Information
                          Traffic Activity : LOW
                          Packets Received Per Second :              20
                          Packets Processed Per Second :              6
                          Connections Created Per Second :            0

                          Avg Packet Size : 128 Bytes
                          Avg TCP Packet Size : 128 Bytes
                          Avg UDP Packet Size : 64 Bytes
```

```
TCP = 76.0 %, UDP = 22.0 % Other: 2.0 %

Packet Size Distribution
Avg-Size    % Share
Bucket[0] (0-128B)     :   64 Bytes    77.47 %
Bucket[1] (128-256B)   :  202 Bytes     7.47 %
Bucket[2] (256-512B)   :  358 Bytes     3.64 %
Bucket[3] (512-1024B)  :  585 Bytes     6.60 %
Bucket[4] (1024+B)     : 1415 Bytes     4.81 %
```

Note: The possible levels of Traffic Activity are: low, medium, high, and critical.

Command Change History

Release	Modification
11.3.1	Added statistics for packets received/processed and connections created per second
8.7.1	**traffic statistics** command introduced

traffic tree

The **traffic tree** command provides detailed information about how often classes and their associated policies are accessed by the PacketShaper classification process, along with rate information for each class. You can also display classes in the traffic tree of the standby partner when using the high availability feature.

```
traffic tree [partner <partner_ip> [<tclass>] [clear]
```

partner *<partner_ip>*	Lists the traffic tree of a standby partner
[*<tclass>*]	The traffic class tree to display, inbound or outbound. If omitted, this defaults to outbound. The class' explicit path must be supplied *only* if the class name itself is not unique.
[clear]	Clears the class and policy hit counts

To view statistics for the entire traffic class tree, use the **traffic tree** command without supplying a specific class name.

Class name	Type	Class hits	Policy hits	Cur rate	1 Min avg	Peak rate
/Inbound	+		n/a	284	366	n/a
Localhost	PE	964	964	0	8	136k
FileMaker		0	n/a	0	0	0

HTTP		0	n/a	0	0	0
POP3		0	n/a	0	0	0
SMTP		0	n/a	0	0	0
SNMP		0	n/a	0	0	0
DNS		0	n/a	0	0	0
NetBIOS-IP		0	n/a	0	0	0
SLP		62	n/a	0	0	487
GRE		0	n/a	0	0	0
ICMP		140	n/a	0	0	108
CiscoDiscovery		2389	n/a	0	142	2391
AppleTalk		1310	n/a	1	3	1058
Default	P I	235	4817	140	213	2298
/Outbound	+		n/a	459	301	n/a
Localhost	PE	968	968	0	3	372k
DHCP		0	n/a	0	0	0
FileMaker		13	n/a	0	0	47.9k
HTTP		229	n/a	226	289	80.1k
EntryPoint		392	n/a	0	0	7338
POP3		12	n/a	0	0	49.0k
SMTP		14	n/a	0	0	12.2k
Telnet		0	n/a	0	0	0
DNS		29	n/a	0	0	150
NetBIOS-IP		75	n/a	0	0	975k
ICMP		129	n/a	0	0	14
Default	P I	27	920	3	9	1.3M

The display shows all of the current traffic classes, with flags that indicate if a class has an associated policy (P), if it is an inheritable class (I), and if it is an exception class (E). A plus sign (+) next to a class represents a partition. This list also shows the number of times a class and its associated policy have been hit. For TCP and UDP traffic classes, PacketShaper counts traffic flows, except for ICMP, for which packets are counted. For non-IP traffic classes, PacketShaper counts packets. The rate statistics for the traffic through the class are also shown.

If a class does not have a policy associated with it, the classification process searches down the tree for a matching sibling that has an inheritable policy. Typically, the default classes such as /Inbound/Default show more policy hits than class hits. The policy hits for the /Inbound/Default class include policy hits for classes without policies listed earlier in the subtree. When a class does not have its own policy, it inherits a policy from a sibling in its subtree.

Command Change History

Release	Modification
11.6.1	**partner** option added

ul add

Add entries to an existing user list. When specifying multiple names, separate each with a space.

```
ul add <user_list> u:<user> | g:<group> [u:<user> | g:<group> ...]
```

where *<user_list>* is an existing user list name, and *<user>* and *<group>* must be specified in the *<domain-name>**<user or group name>* format. To differentiate between user and group names, you must precede each name with **u:** or **g:**.

Example:

```
ul add list1 g:cal\group-marketing u:cal\john.smith u:cal\peter.hanson
```

Command Change History

Release	Modification
9.2.2	Support added for user groups; command syntax modified
9.2.1	Command introduced

ul delete

Remove one or more users or groups from an existing user list.

```
ul delete <user_list> u:<user> | g:<group> [u:<user> | g:<group> ...]
```

where *<user_list>* is an existing user list name, and *<user>* and *<group>* must be specified in the *<domain-name>*\|*<user or group name>* format.

Command Change History

Release	Modification
9.2.2	Support added for user groups; command syntax modified
9.2.1	Command introduced

ul new

Create a user list by defining a unique name and specifying the user and group names that should be included in the list. When specifying multiple names, separate each with a space.

```
ul new <user_list> u:<user> | g:<group> [u:<user> | g:<group> ...]
```

where *<user_list>* is a descriptive name, up to 127 characters; the slash (/) and backslash (\) characters may not be used. The *<user>* and *<group>* must be specified in the *<domain-name>*\|*<user or group name>* format. To differentiate between user and group names, you must precede each name with **u:** or **g:**.

To add entries to the user list after it's created, use the ul add command.

Examples:

```
ul new list2 u:ny\pharrison u:ny\speters g:ny\group-sales
```

```
ul new list3
```

> **Note:** The **ul new** command does not validate the existence of the user and group names.

Command Change History

Release	Modification
9.2.2	Support added for user groups; command syntax modified
9.2.1	Command introduced

ul override

For PolicyCenter/ Units in shared mode only

Override an inherited user list by creating a local copy of the list.

```
ul override <list-name>
```

where *<user_list>* is an existing user list name. You must make a local copy of an inherited user list before you can change the user list on the child configuration.

ul reformat

This command ensures v9.2.1 user lists are formatted properly after upgrading to v9.2.2 or higher.

```
ul reformat upgrade
```

The **upgrade** option adds **u:** before each user name in the list because a prefix is necessary to distinguish user names from group names in v9.2.2+. It performs the operation on all user lists at once.

> **Note:** If you fail to use the upgrade command, user list classification will still work properly. However, you will see an "invalid entries in the list" error if you ever try to modify the list, and you will need to manually add the **u:** prefix on each user name (or run the **ul reformat upgrade** command to fix all user lists).

Command Change History

Release	Modification
9.2.2	New command

ul rm

Remove a user list from the configuration.

```
ul rm <user_list>
```

User lists cannot be removed if they are currently being used in a class matching rule.

Command Change History

Release	Modification
9.2.1	Command introduced

ul show

Display a list of all defined user lists or show the details of a specific user list.

```
ul show [<list_name>]
```

To show all user lists, as well as the user and group names in each list:

```
ul show *
```

User and group names are listed alphabetically.

Command Change History

Release	Modification
9.2.2	Support added for user groups
9.2.1	Command introduced

uptime

Determine how long the unit has been up and running. It measures the time since the unit was booted, either from a power-up or a software reset.

```
uptime
```

version

Display the PacketShaper software version that is currently running, build number (release ID of the active image), model, and serial number. Use the **verbose** option to list the part number, the inside and outside MAC addresses, and installed license keys.

```
version [verbose]
```

The output of **version verbose** will show *(Not applicable)* next to any key that you haven't purchased.

watch add

Add a router to the watch mode router list (the routers whose traffic is being monitored).

```
watch add <name> <MAC address>
```

where

<name>	Description of router; up to 32 characters (no spaces are allowed, the only special characters allowed are colon, dash, underline, and period.)
<MAC address>	MAC address of the router, for example *08:00:20:c0:56:a6*

Example:

```
watch add router1 08:00:20:c0:56:a6
```

A PacketShaper in watch mode can monitor traffic from up to 256 routers.

Command Change History

Release	Modification
11.1.1	Removed *<IP address>* option

watch delete

Delete a router from the watch mode router list.

```
watch remove <name>
```

where *<name>* is the name that was defined when the router was added. To see a list of defined router names, use the watch show command.

Example:

```
watch remove router1
```

watch show

Display the watch mode configuration. The output lists the current management port and configured watch mode routers.

```
watch show
```

Example:

```
watch show
```

```
Watch Mode Status: Enabled
 Management Port: MGMT

Name              MAC Address
-----------------------------------------------------------
router6           00:10:7b:3c:30:39
router5           08:00:20:c0:56:a6
router4           00:03:e3:6b:46:c2
router3           00:03:e3:6b:46:c2
router2           01:02:03:04:05:06
router10           00:60:fb:60:1f:16
```

The following information is displayed in this screen:

Field	Description
Watch Mode Status	Indicates whether watch mode is currently enabled or disabled
Management Port	Indicates which port PacketShaper has determined will be used to manage the unit. The management port is not user-definable. PacketShaper decides which port to use for management access by checking which ports are connected.
Name	The list of configured routers.
MAC Address	Use the watch add command to configure routers and the watch delete command to remove them from the list.

Bytes Transmitted Graph

The Bytes Transmitted graph compares the number of transmitted bytes to the number of retransmitted bytes for the link, a class, or several classes.

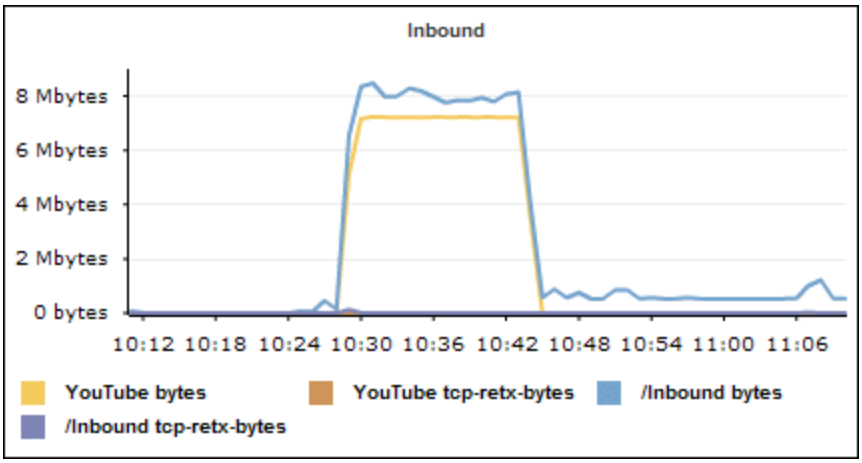

Connection Retransmissions Graph

The Connection Retransmissions graph compares the TCP retransmit rate (the percentage of packets that were retransmitted) with the TCP toss rate (the percentage of retransmitted packets that were tossed) of a link, a class, or several classes. This graph can help identify poorly tuned servers whose retransmissions are wasting bandwidth.

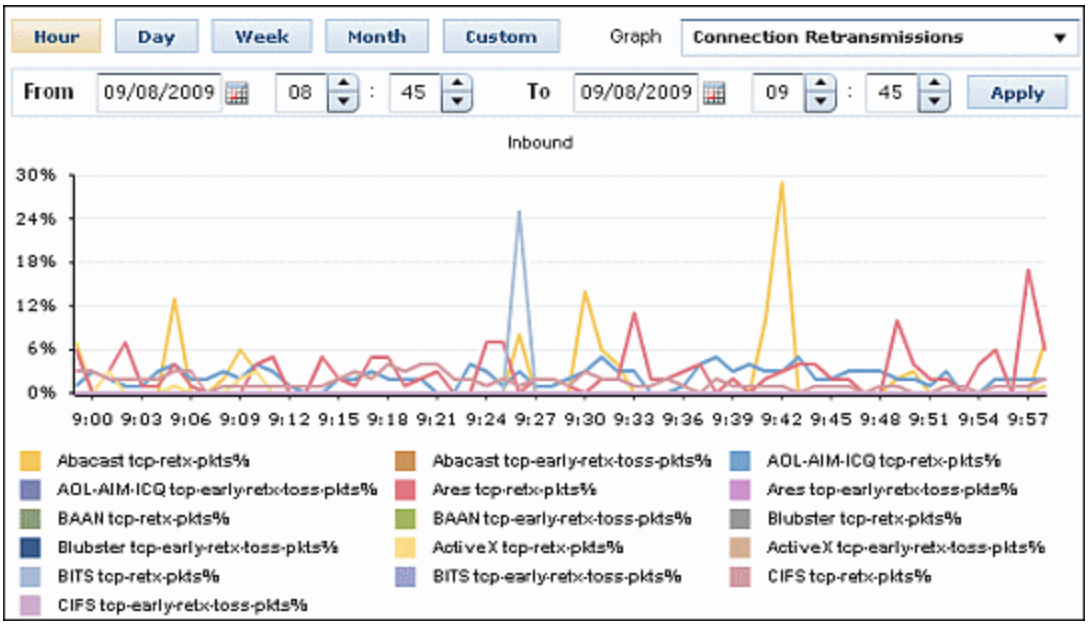

Guaranteed Rate Failures Graph

The Guaranteed Rate Failures graph shows a history of the number of times PacketShaper was unable to provide the bandwidth guaranteed by the class' policy. If the selected class does not have an associated guaranteed rate, the graph shows zeroes.

If this graph shows an unacceptable increase in rate failures, ask the following questions:

- Is the rate guarantee inappropriate?
- Am I having an unexpected number of flows?
- Is a partition or a lower guarantee workable?
- Is my policy's admission control set adequately?

Network Efficiency Graph

The Network Efficiency graph shows the amount of wasted TCP traffic by displaying the percentage of bytes that are *not* retransmits. This graph displays the measurement engine variable tcp-data-bytes as a percentage, which is computed as:

```
(bytes - tcp-retx-bytes) / bytes
```

> **Note:** If the selected class' traffic is not TCP, then the graph indicates 100%.

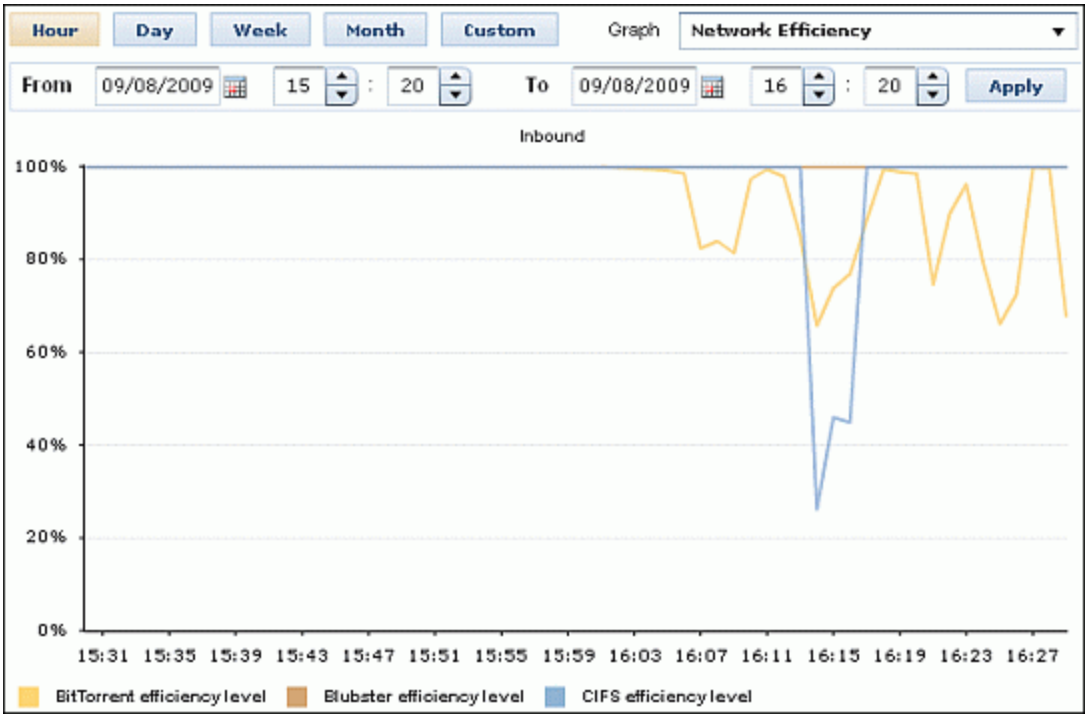

The Network Efficiency graph can help you spot bad trends. Tosses and retransmitted packets lower the efficiency percentage. An efficient network, displayed as minor fluctuations from 100%, needs very little intervention. A lower percentage indicates that more of your network capacity is devoted to retransmitting packets. In this case, you may want to reevaluate your current partitions and policies. See Modify a Policy.

Normalized Network Delay Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Normalized Network Delay graph shows the transaction delay in the network, normalized by transaction size. It shows how long it takes to send 1KB of data. This graph allows an accurate comparison of response-time data for different applications or servers. Without normalizing the delay, response times vary depending on the size of the transaction. This graph eliminates size as a factor of network delay.

> **Note:** The Normalized Network Delay graph is based on the normalized-network-delay-avg variable.

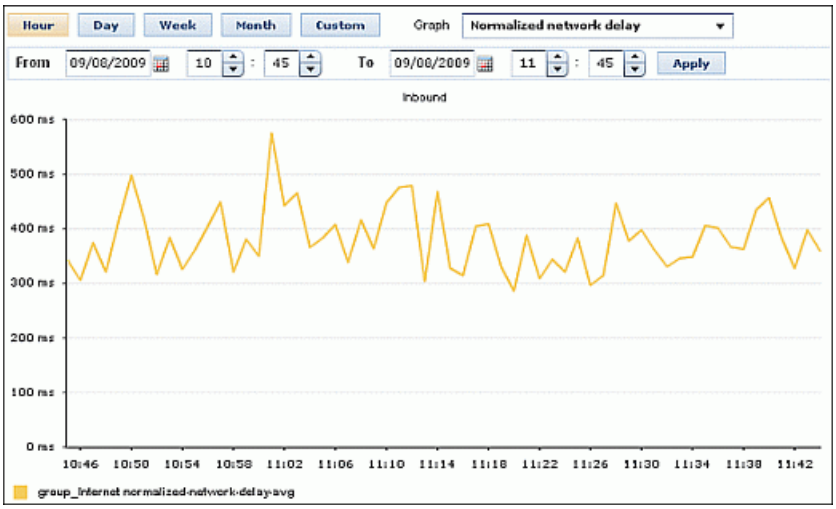

Packet Round Trip Time Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Packet Round Trip Time graph shows the history of a transaction's round-trip time (RTT). RTT is the average number of milliseconds spent in transit when a client and server exchange the SYN (synchronize sequence numbers flag) and its corresponding ACK (acknowledge flag). A transaction involving a large amount of data requires the data to be divided into multiple packets. Whereas a transaction's network delay reflects the total transit time for all required packets, the RTT reflects the time for a single packet to make its way from client to server and another packet to make the return trip.

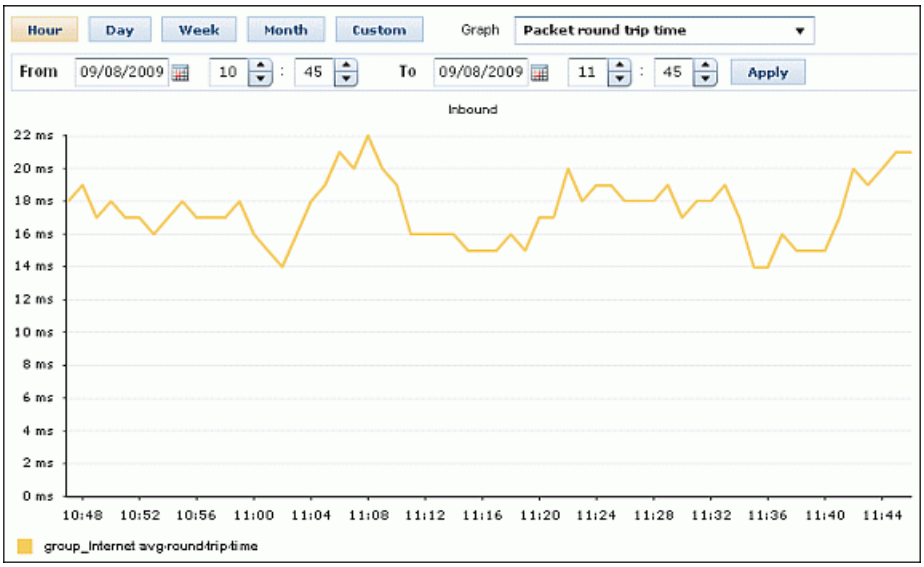

The RTT graph provides a good way to track the history of network speed, independent of the size of historical transactions. The RTT from a given client to a given server only varies with network speed; the size of the transaction is irrelevant.

Packets Transmitted Graph

The Packets Transmitted graph compares the number of transmitted packets to the number of retransmitted packets (excluding tossed packets) for a the link, a class, or several classes.

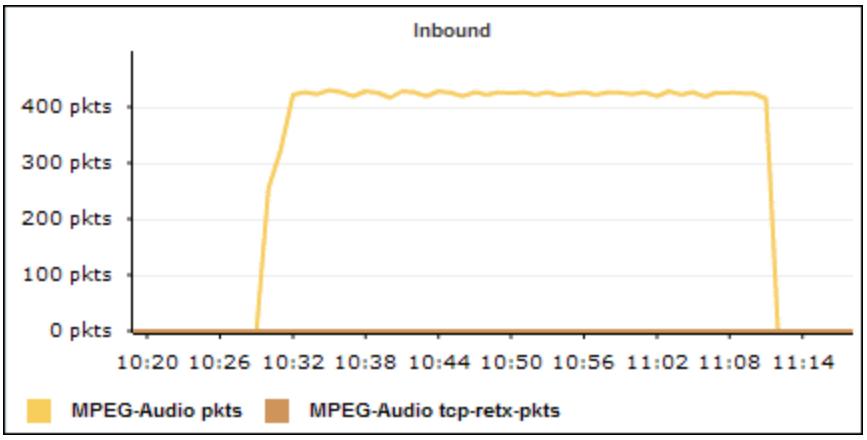

TCP Connections Initiated Graph

The TCP Connections Initiated graph shows the number of TCP connections started in the specified time period.

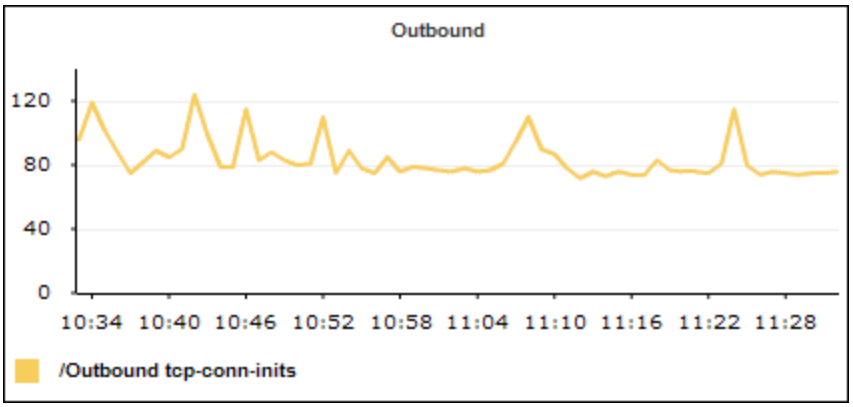

TCP Health Graph

The TCP Health graph gives you a comprehensive picture of TCP connections for a link, class, or several classes. It compares the number of TCP connections that were started, aborted, ignored by the server (that is, the server never responded), and refused by the server during the specified time period.

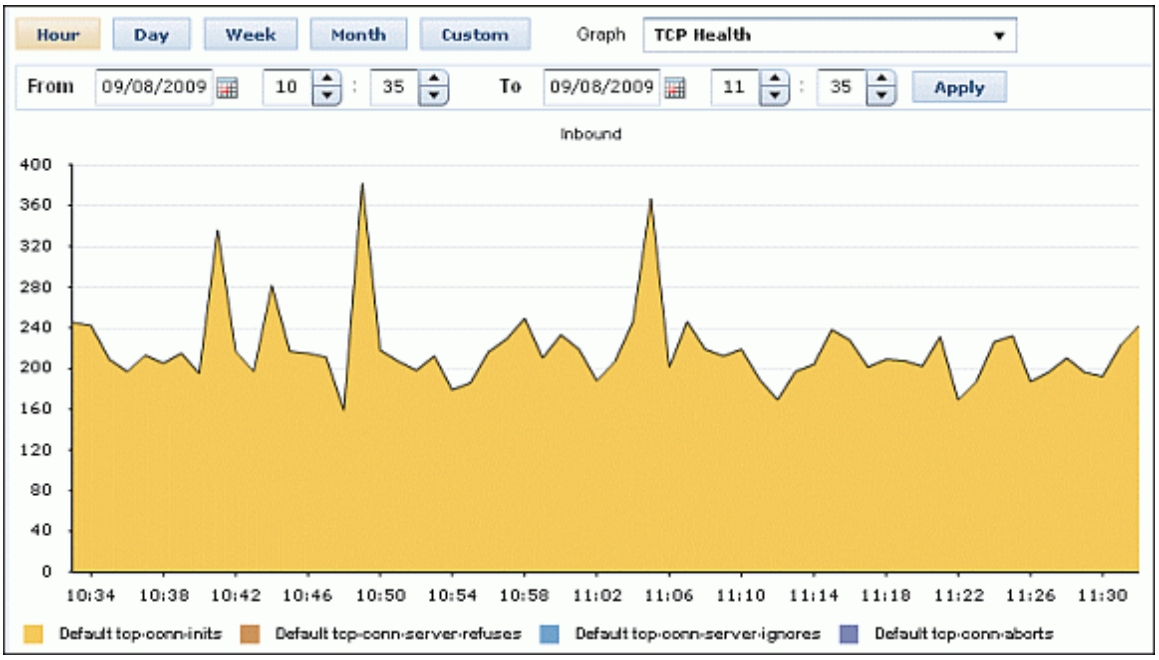

For example, suppose you created classes for each branch office and one of your offices complained about a server problem (such as slow speed, or many disconnects, or hung systems). By comparing TCP Health graphs of the various branch office classes, you can figure out what is causing the server problem at the one branch office.

For details on each of the measurement variables used in this graph (tcp-conn-inits, tcp-conn-aborts, tcp-conn-server-ignores, tcp-conn-server-refuses), see Measurement Variables.

Transaction Delay Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Transaction Delay graph shows a timeline of the average response times in milliseconds of one or more traffic classes. Three colored lines track the total, network, and server delays, so you can determine the source of any significant delays. The horizontal line represents the Total Delay Threshold as a reference point to determine the quality of the response times.

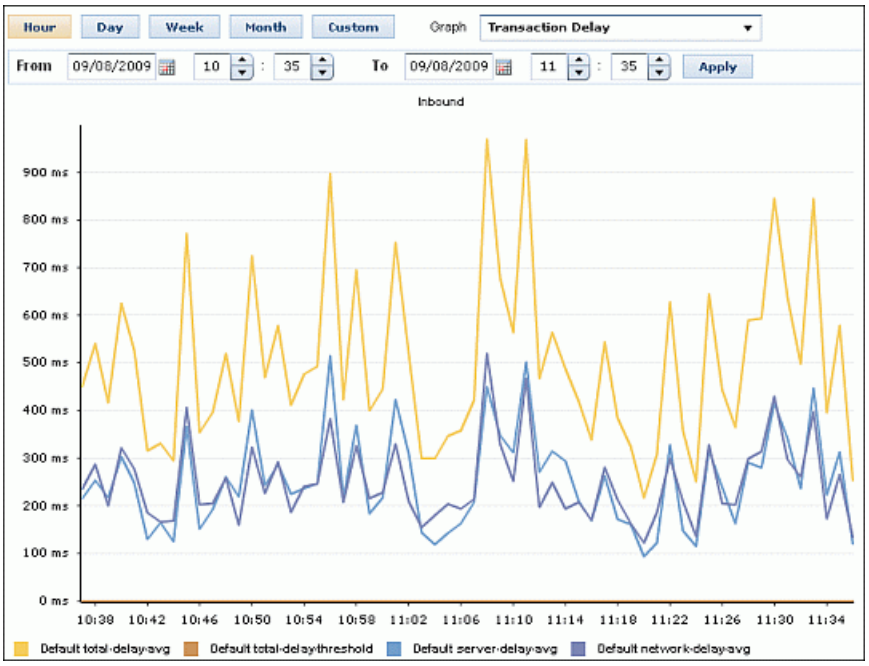

Average Utilization Graph

The Class Utilization, Average graph shows a history of the class' average bandwidth consumption in bits per second. This graph answers questions such as, "How much bandwidth does my FTP traffic class typically take?"

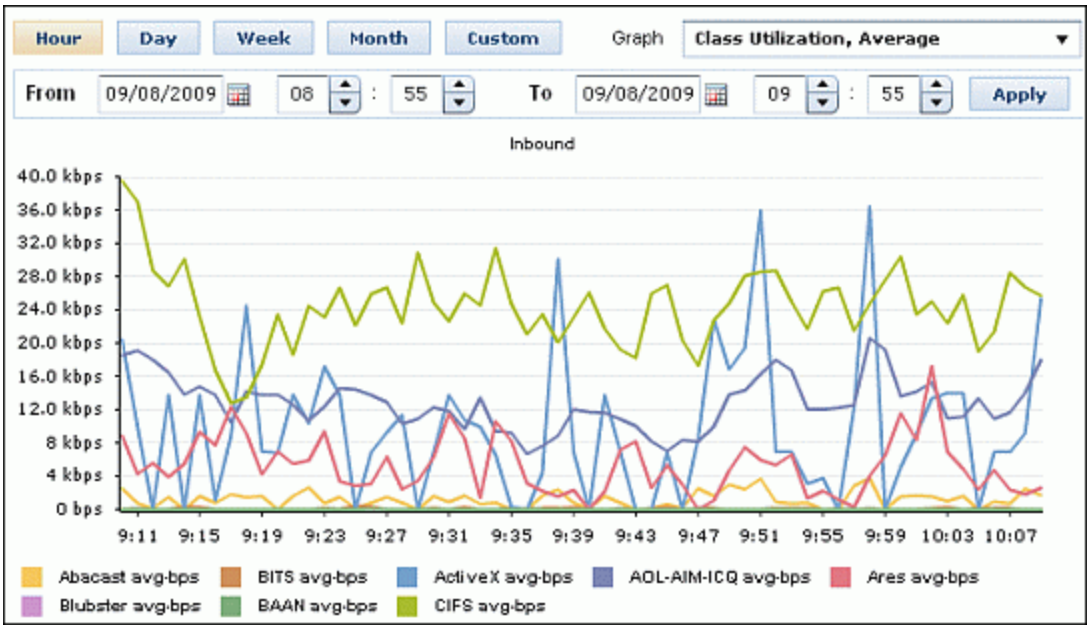

Average bandwidth (without peaks) may show ample capacity, especially when longer time units are measured. By checking the peaks, you can see if a traffic class is frequently hitting a capacity limit. See Peak Utilization Graph.

Peak Utilization Graph

The Peak Utilization graph displays the peak bandwidth consumption, in bits per second, of one or more traffic classes. PacketShaper determines the peak rate by looking at the rate recorded for the busiest one-second sub-interval (that is, the sub-interval that had the highest rate).

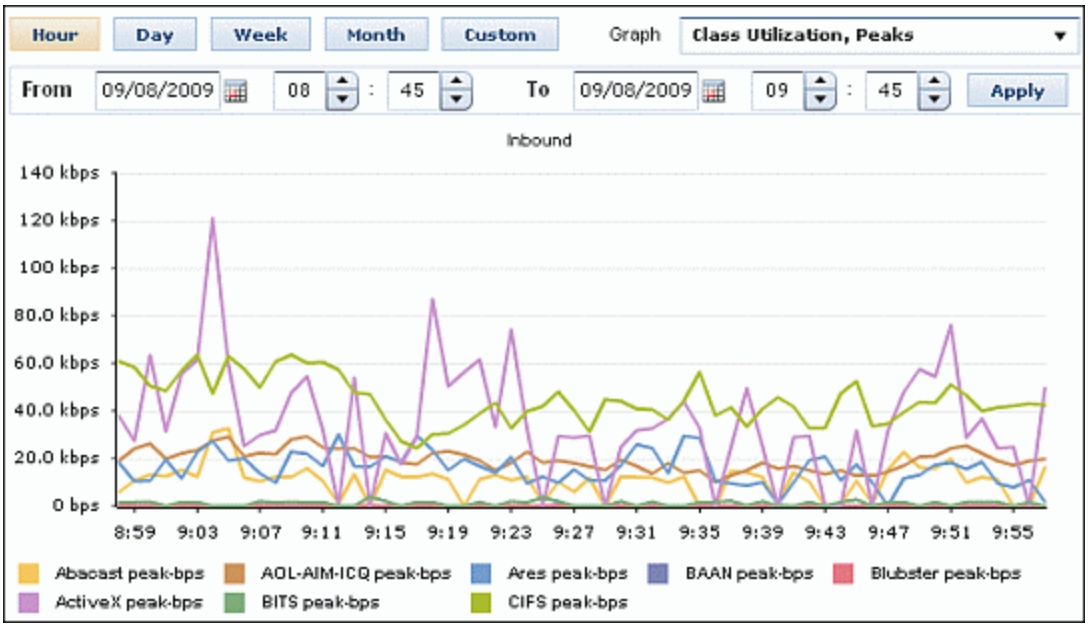

Average Transaction Size Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Average Transaction Size graph shows the average number of bytes per transaction of a TCP-based application.

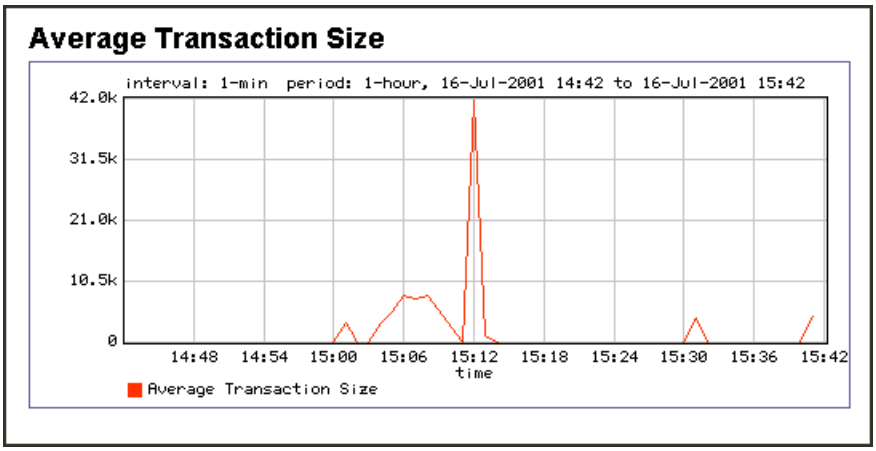

Bytes Transmitted Graph

The Bytes Transmitted graph compares the number of transmitted bytes to the number of retransmitted bytes for a specific link, partition, or class.

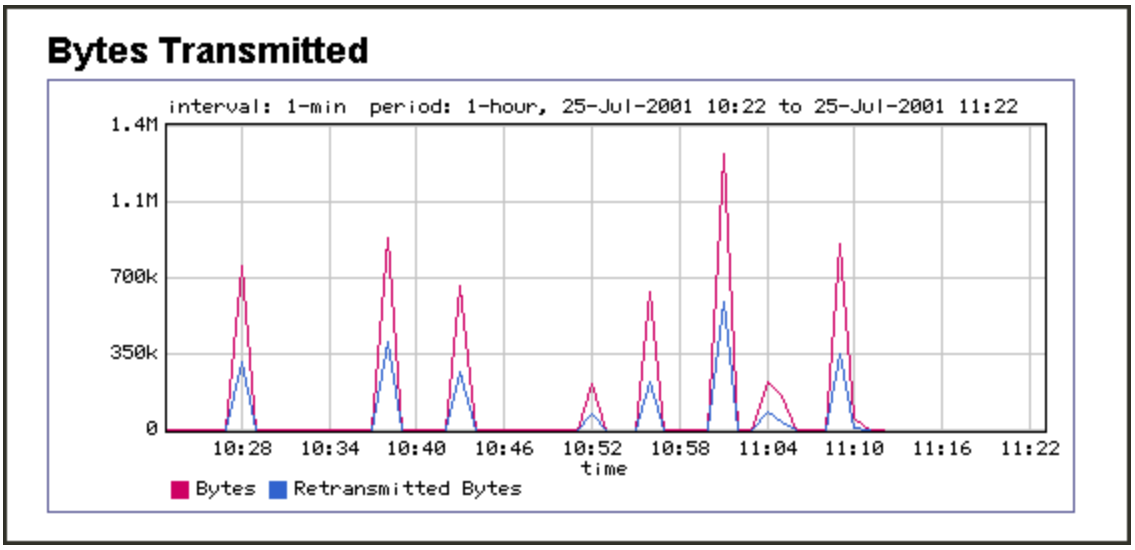

Class Utilization with Peaks Graph

The Class Utilization with Peaks graph displays a traffic class' average and peak bandwidth consumption, in bits per second, over time. PacketShaper determines the peak rate by looking at the rate recorded for the busiest one-second sub-interval (that is, the sub-interval that had the highest rate).

Note: Because the average rate is weighted over five seconds, you might find that sometimes this rate is higher than the peak rate for a given interval. You might occasionally see artificial spikes in the peaks and averages; this happens when the collection and processing of data is paused because higher priority tasks are executing.

Tip: A quick way to display this graph is to go to the Top Ten window and click the graph icon next to the class you want to graph.

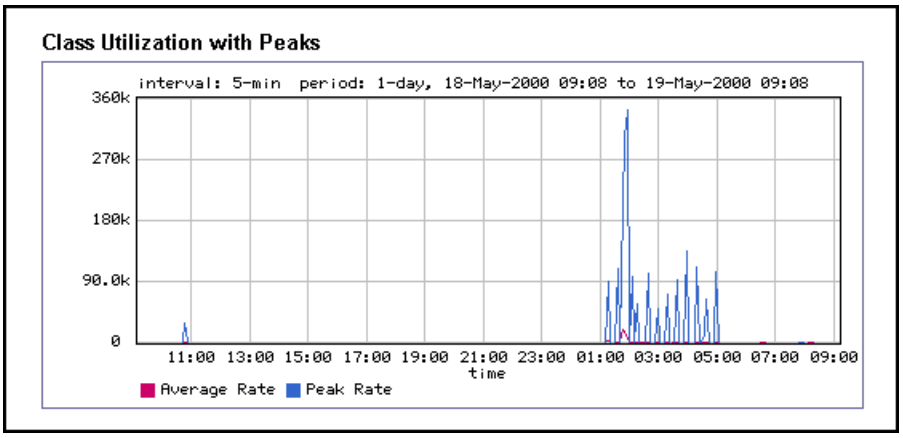

Class Utilization Graph

The Class Utilization graph shows a history of the class' average bandwidth consumption in bits per second. This graph answers questions such as, "How much bandwidth does my FTP traffic class typically take?"

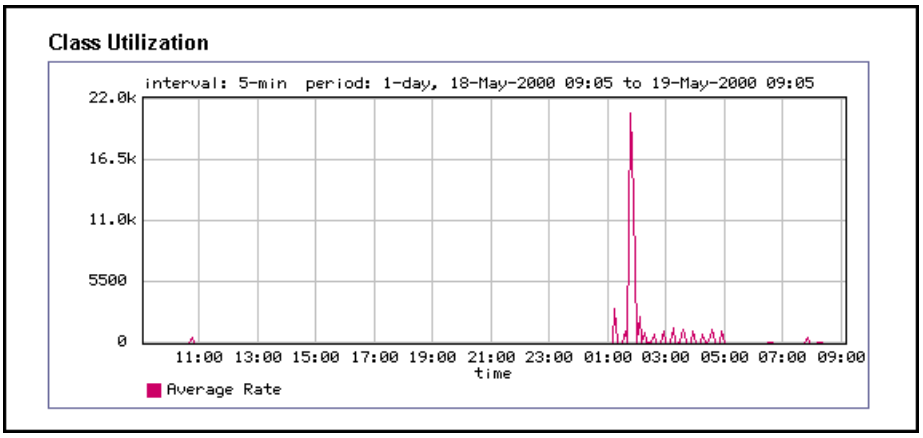

Although the average bandwidth utilization over time is also displayed in the Class Utilization graph, the same figures may look different in the Class Utilization with Peaks graph. This is because the graph's scale changes if there is a sizeable difference between the average and peak figures.

Average bandwidth (without peaks) may show ample capacity, especially when longer time units are measured. By checking the peaks, you can see if a traffic class is frequently hitting a capacity limit.

Connection Retransmissions Graph

The Connection Retransmissions graph compares the TCP retransmit rate (the percentage of packets that were retransmitted) with the TCP toss rate (the percentage of retransmitted packets that were tossed) of a link, partition, or class. This graph can help identify poorly tuned servers whose retransmissions are wasting bandwidth.

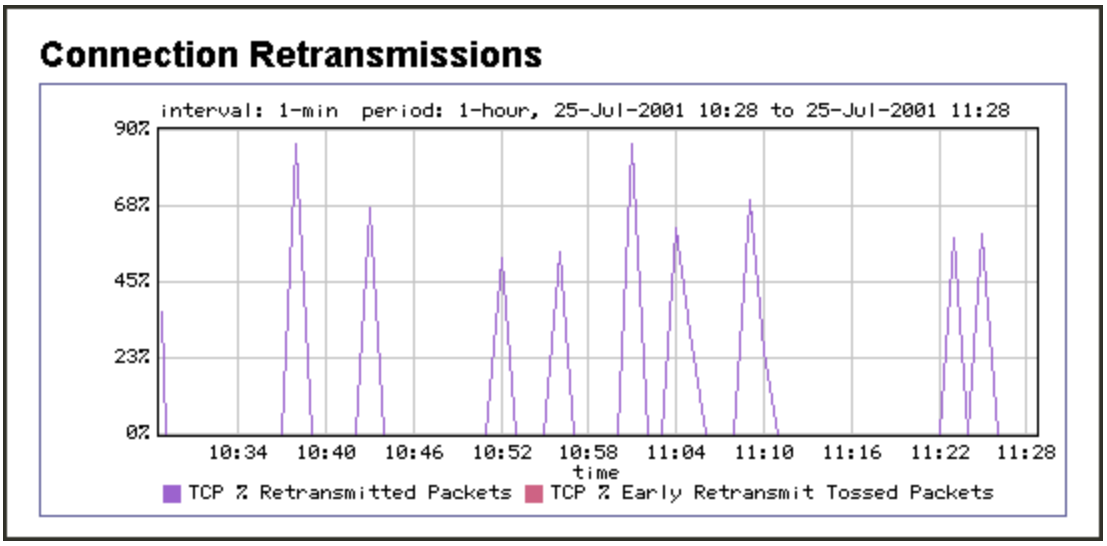

Dynamic Partition Usage

The Dynamic Partition Usage graph provides three statistics pertaining to a particular dynamic partition:

- The number of current users (includes active and idle subpartitions)
- The number of subpartitions PacketShaper attempted to create after the partition's cap was reached.
- Number of dynamic subpartitions that PacketShaper attempted to create but could not because the number of partitions on the unit had reached its limit. The maximum number of partitions on a unit depends on the model.

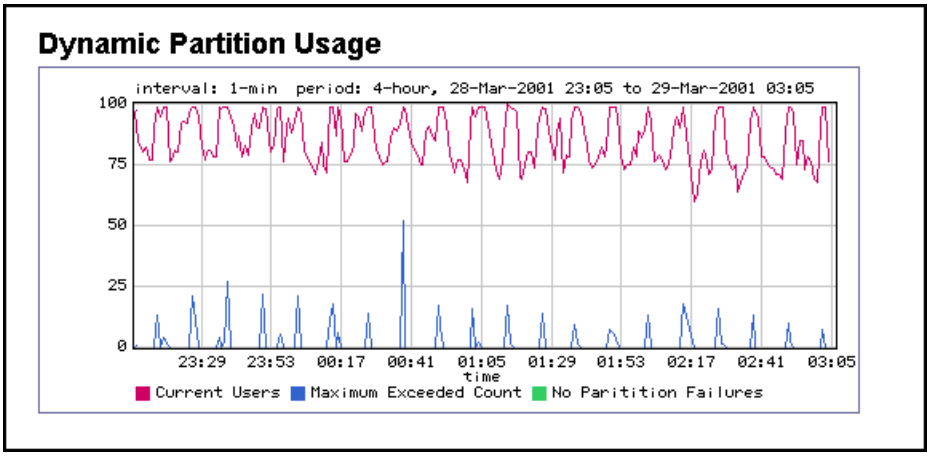

Guaranteed Rate Failures Graph

The Guaranteed Rate Failures graph shows a history of the number of times PacketShaper was unable to provide the bandwidth guaranteed by the class' policy. If the selected class does not have an associated guaranteed rate, the graph shows zeroes.

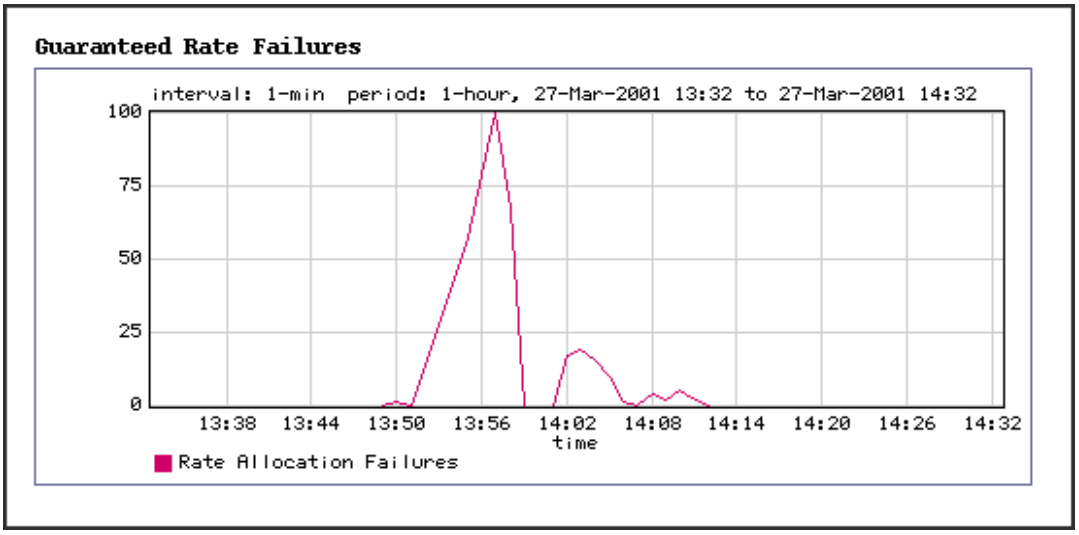

If this graph shows an unacceptable increase in rate failures, ask the following questions:

- Is the rate guarantee inappropriate?
- Am I having an unexpected number of flows?
- Is a partition or a lower guarantee workable?
- Is my policy's admission control set adequately?

Link Utilization with Peaks Graph

The Link Utilization with Peaks graph shows the link's average and peak bandwidth usage in bits per second. PacketShaper determines the peak rate by looking at the rate recorded for the busiest one-second sub-interval (that is, the sub-interval that had the highest rate).

> **Note:** Because the average rate is weighted over five seconds, you might find that sometimes this rate is higher than the peak rate for a given interval. You might occasionally see artificial spikes in the peaks and averages; this happens when the collection and processing of data is paused because higher priority tasks are executing.

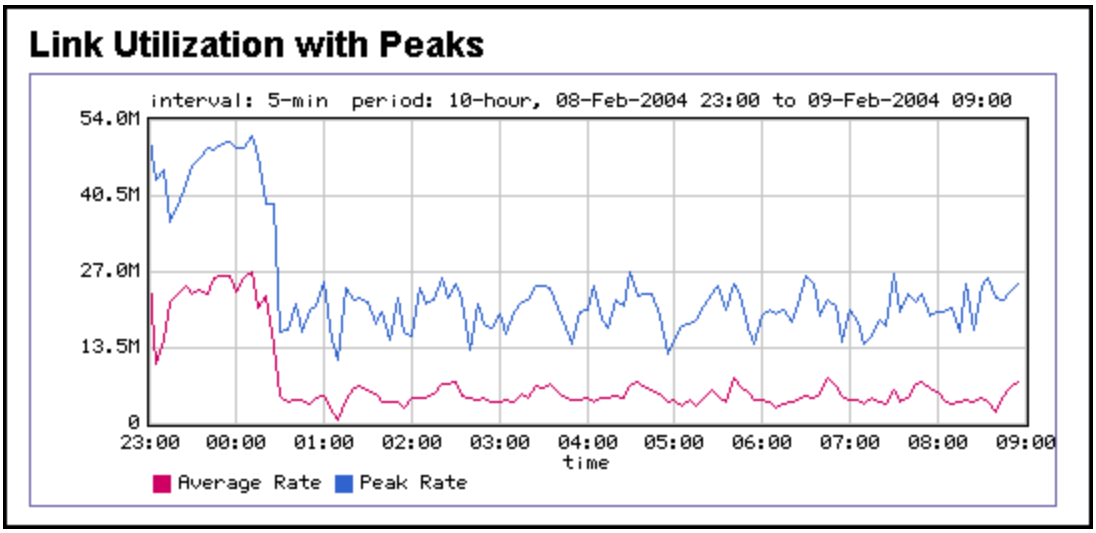

If you choose to display the link size when creating the graph, a horizontal line indicates the capacity of the link. When the link size is displayed (as shown below), the graph can answer questions such as, "How frequently is my link size insufficient? What are my average capacity needs?"

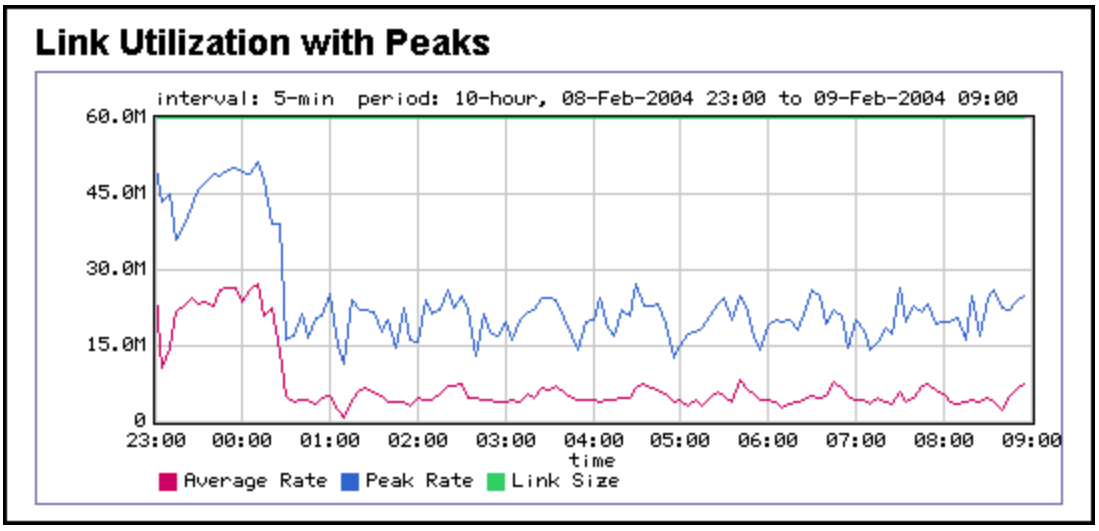

Link Utilization Graph

The Link Utilization graph shows the link's average bandwidth usage in bits per second. If you choose to display the link size when creating the graph, a horizontal line indicates the capacity of the link, as shown below.

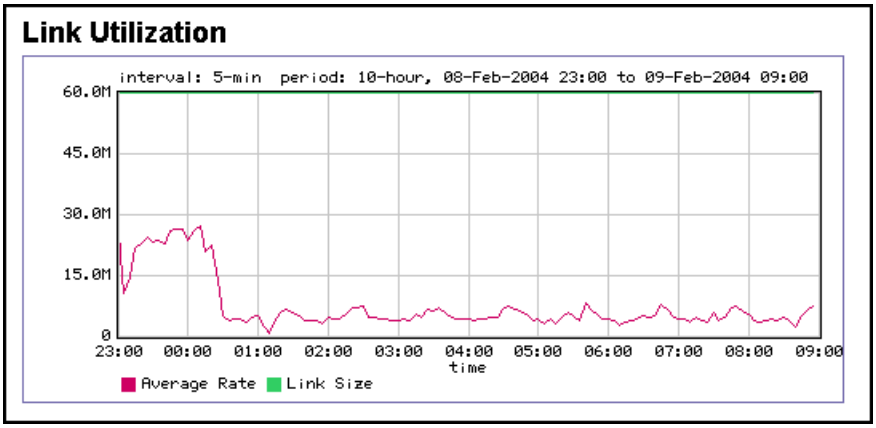

This graph answers questions such as "Does the link usage vary a lot? What are my average capacity needs?"

> **Note:** To get the entire picture, you should also view the Link Utilization with Peaks graph (using sufficiently small time intervals).

Network Delay Distribution

> **Note:** This graph type is available only if RTM is available for the selected object.

The Network Delay Distribution graph displays the number of transactions whose network delay falls into each of the 14 transaction-time buckets described under the Transaction Delay Distribution graph. The median delay is indicated on the bottom of the graph.

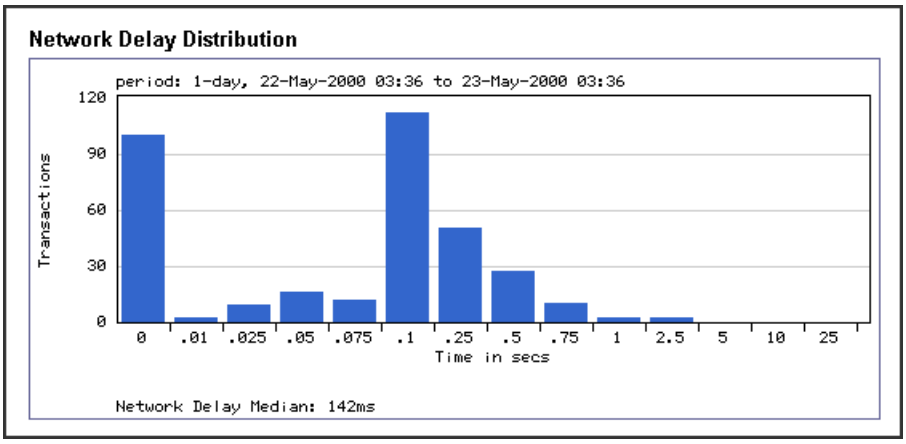

Network Delay Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Network Delay graph shows the average response times in milliseconds of a traffic class over time. This graph shows only the portion of the transaction time that is attributable to the network, enabling you to analyze network delay.

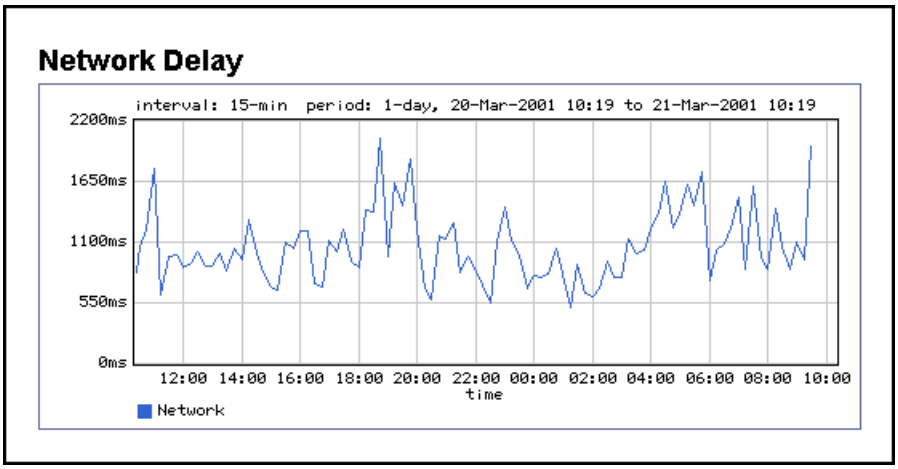

Network Efficiency Graph

The Network Efficiency graph shows the amount of wasted TCP traffic by displaying the percentage of bytes that are *not* retransmits. This graph displays the measurement engine variable tcp-data-bytes as a percentage, which is computed as:

```
(bytes - tcp-retx-bytes) / bytes
```

> **Note:** If the selected class' traffic is not TCP, then the graph indicates 100%.

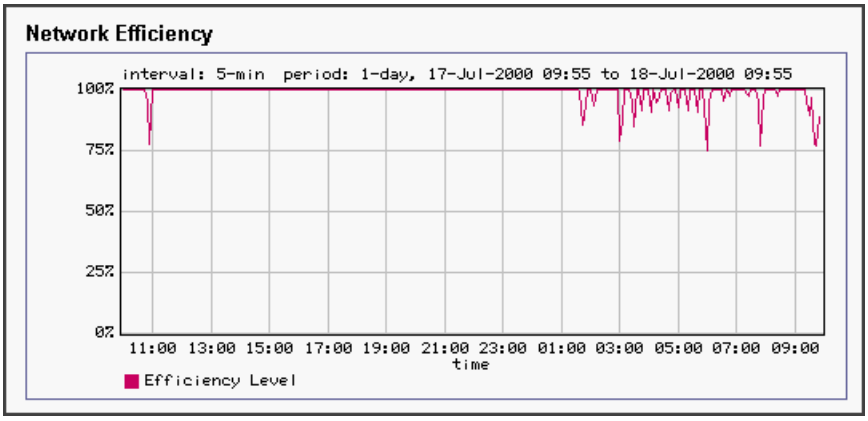

The Network Efficiency graph can help you spot bad trends. Tosses and retransmitted packets lower the efficiency percentage. An efficient network, displayed as minor fluctuations from 100%, needs very little intervention. A lower percentage indicates that more of your network capacity is devoted to retransmitting packets. In this case, you may want to reevaluate your current partitions and policies. For more information, see Modify a Partition and Modify a Policy.

Normalized Network Delay Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Normalized Network Delay graph shows the transaction delay in the network, normalized by transaction size. It shows how long it takes to send 1KB of data. This graph allows an accurate comparison of response-time data for different applications or servers. Without normalizing the delay, response times vary depending on the size of the transaction. This graph eliminates size as a factor of network delay.

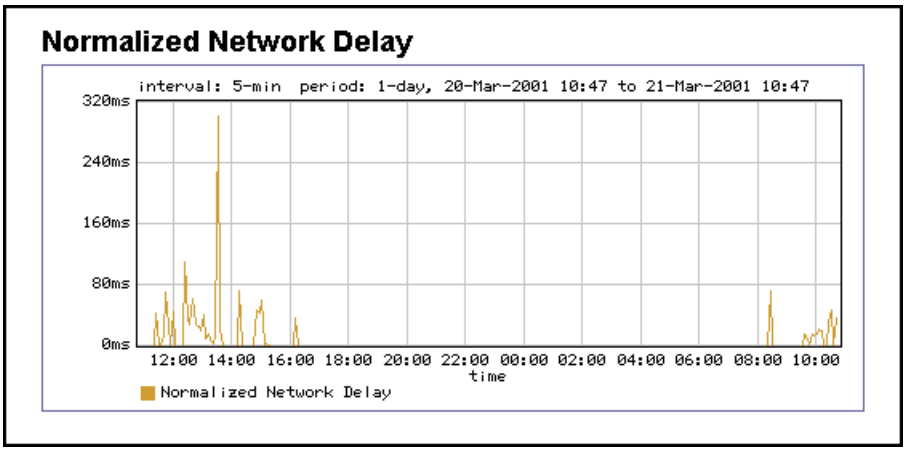

Packet Round-Trip Time Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Packet Round-Trip Time graph shows the history of a transaction's round-trip time (RTT). RTT is the average number of milliseconds spent in transit when a client and server exchange the SYN (synchronize sequence numbers flag) and its corresponding ACK (acknowledge flag). A transaction involving a large amount of data requires the data to be divided into multiple packets. Whereas a transaction's network delay reflects the total transit time for all required packets, the RTT reflects the time for a single packet to make its way from client to server and another packet to make the return trip.

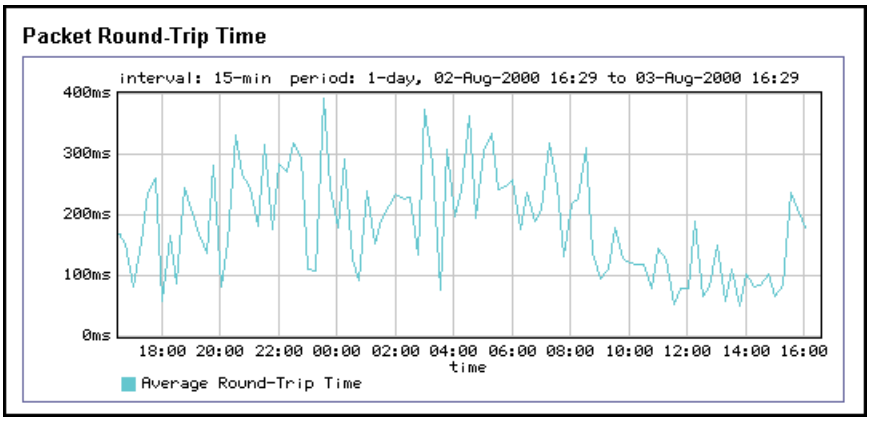

The RTT graph provides a good way to track the history of network speed, independent of the size of historical transactions. The RTT from a given client to a given server only varies with network speed; the size of the transaction is irrelevant.

Packet Size Distribution Graph

The Packet Size Distribution graph is a histogram of packets received on the Inbound or Outbound link, in seven different size buckets. The buckets of packet sizes (in bytes) include: [0-63], [64-127], [128-255], [256-511], [512-1023], [1024-1517], [>=1518]. In the example below, about 340,000 packets were received that were in the "127" bucket (that is, the size range of 64-127 bytes).

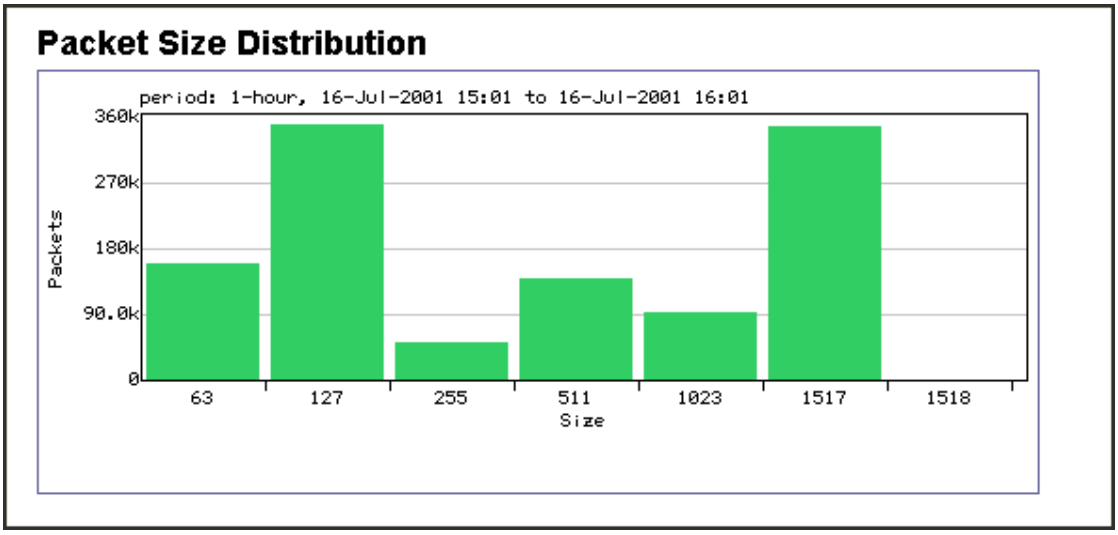

Packets Transmitted Graph

The Packets Transmitted graph compares the number of transmitted packets to the number of retransmitted packets (excluding tossed packets) for a specific link, partition, or class.

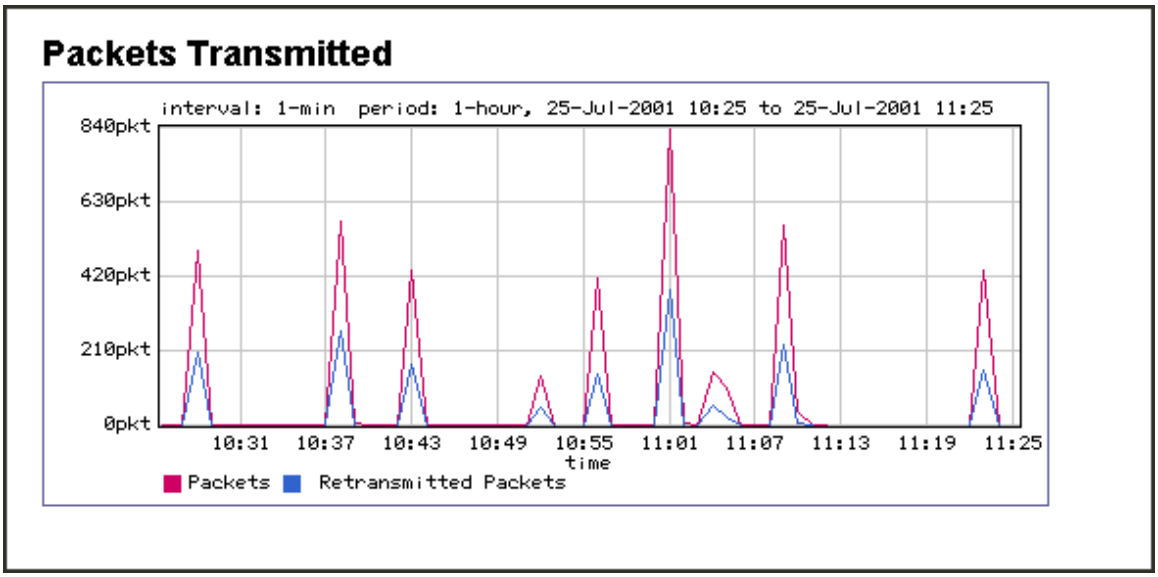

Partition Utilization with Peaks Graph

The Partition Utilization with Peaks graph shows a partition's average and peak bandwidth usage in bits per second. PacketShaper determines the peak rate by looking at the rate recorded for the busiest one-second sub-interval (that is, the sub-interval that had the highest rate).

> **Note:** Because the average rate is weighted over five seconds, you might find that sometimes this rate is higher than the peak rate for a given interval. You might occasionally see artificial spikes in the peaks and averages; this happens when the collection and processing of data is paused because higher priority tasks are executing.

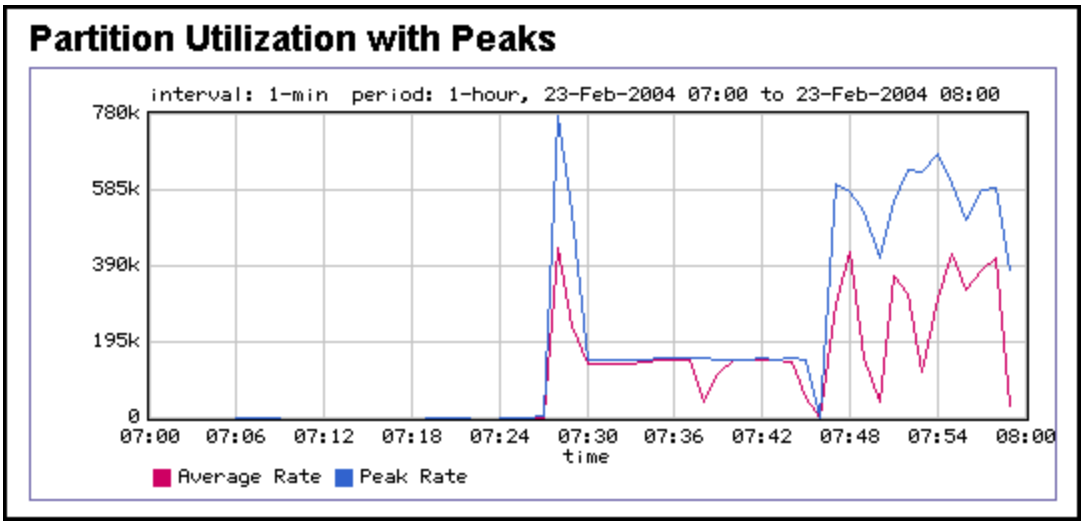

If you choose to display the partition size when creating the graph, horizontal lines indicate the partition's minimum and burstable size limits (as shown below). This graph can answer questions such as, "I reserved 150 Kbps for this partition is all that bandwidth really needed?" In the graph below, the partition was changed from non-burstable to burstable during the time period that was graphed. It is easy to see that excess bandwidth was utilized when the partition was made burstable.

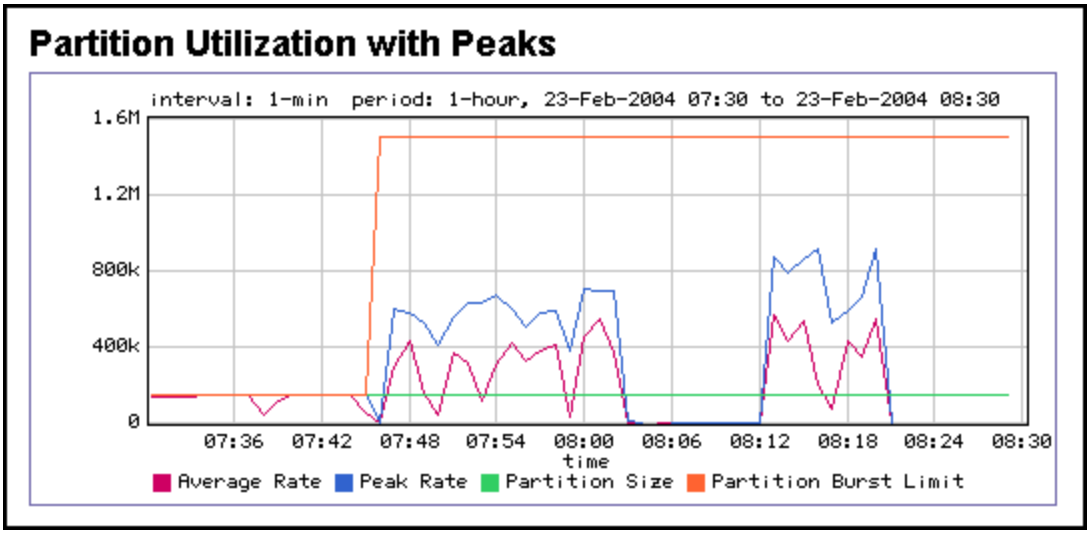

Partition Utilization Graph

The Partition Utilization graph shows a partition's average bandwidth usage in bits per second. If you choose to display the partition size when creating the graph, horizontal lines indicate the partition's minimum and burstable size limits (as shown below). In the graph below, the partition was changed from non-burstable to burstable during the time period that was graphed. It is easy to see that excess bandwidth was utilized when the partition was made burstable.

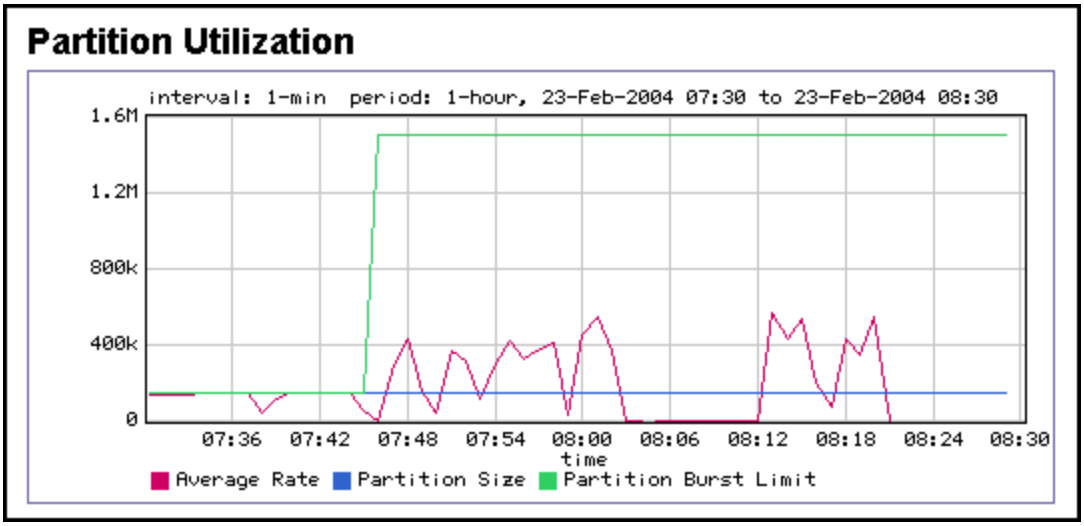

Note: If the selected traffic class does not have a partition, PacketShaper references the partition of the nearest ancestor traffic class.

Server Delay Distribution Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Server Delay Distribution graph displays the number of transactions whose server delay falls into each of the 14 transaction-time buckets described under the Transaction Delay Graph. The median delay is indicated on the bottom of the graph.

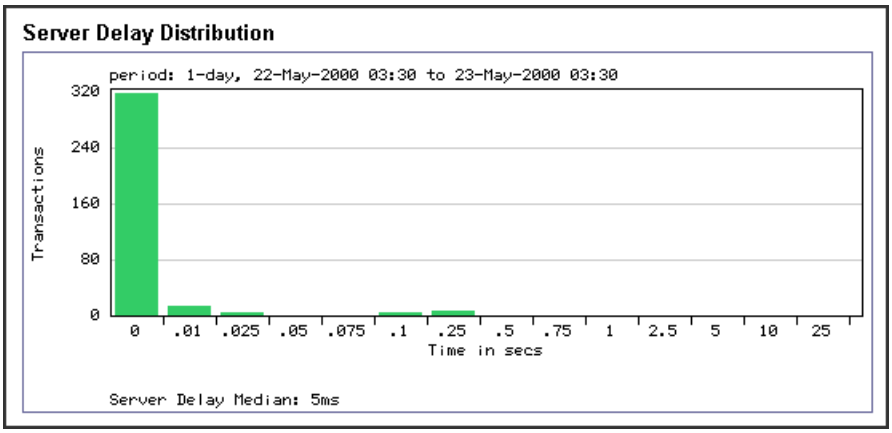

Note that the counts of each type of delay for any one transaction time slot aren't necessarily associated with the same transactions. For example, one transaction might have:

- .025-second server delay (the 3rd slot)
- .475-second network delay (7th slot)
- .5-second total delay (8th slot)

Server Delay Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Server Delay graph shows the average response times in milliseconds of a traffic class over time. This graph shows only the portion of the transaction time that is attributable to the server, enabling you to analyze server delay.

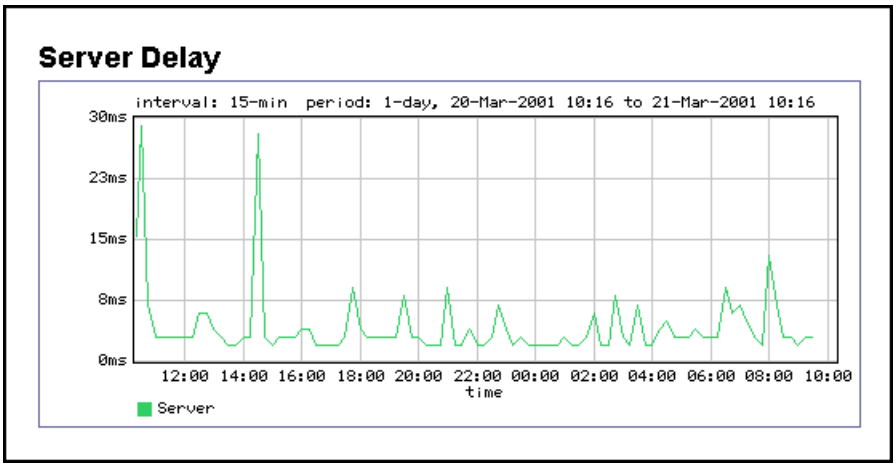

Service Level Compliance Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Service Level Compliance graph is available for traffic classes with response-time measurements and thresholds. This graph shows two lines, one representing the required percentage of good transactions and the other representing the actual percentage of good transactions. The first line is always flat. The second line varies, depending on performance.

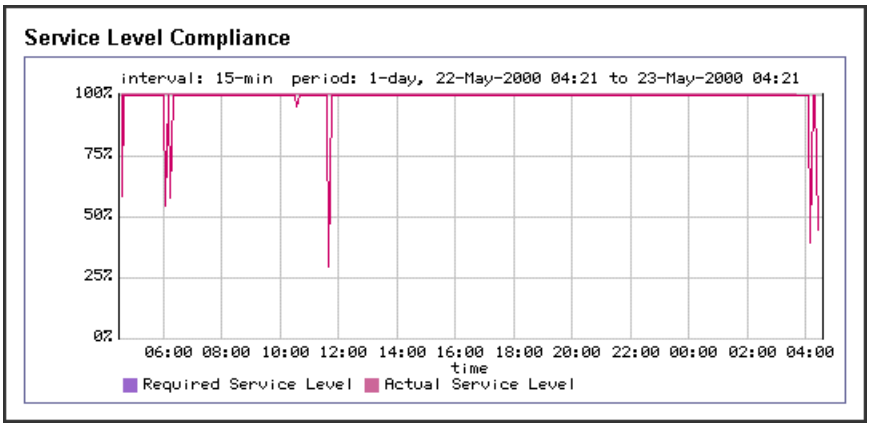

The Service Level Compliance graph shows how a class measures up to its own performance standard. For example, suppose a class has a Total Delay Threshold set to 400 ms and a Service Level Threshold set to 90 percent. For that class to be performing well, at least 90 percent of the class' transaction delays in each one-minute interval should be under 400 ms.

TCP Connections Initiated Graph

The TCP Connections Initiated graph shows the number of TCP connections started in the specified time period.

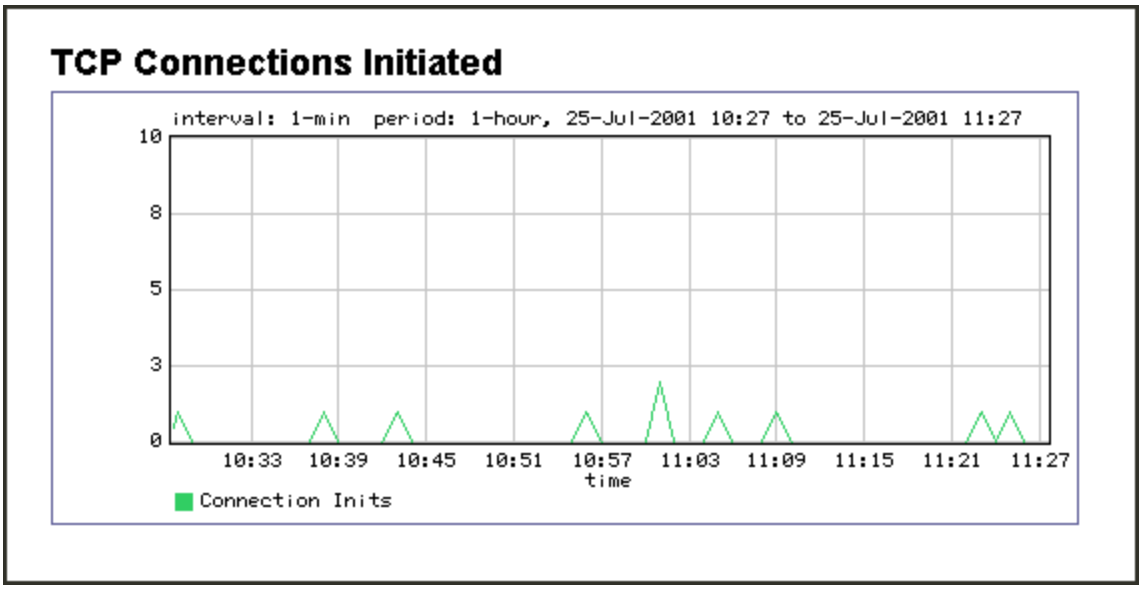

TCP Health Graph

The TCP Health graph gives you a comprehensive picture of TCP connections for a particular link, partition, or class. It compares the number of TCP connections that were started, aborted, ignored by the server (that is, the server never responded), and refused by the server during the specified time period.

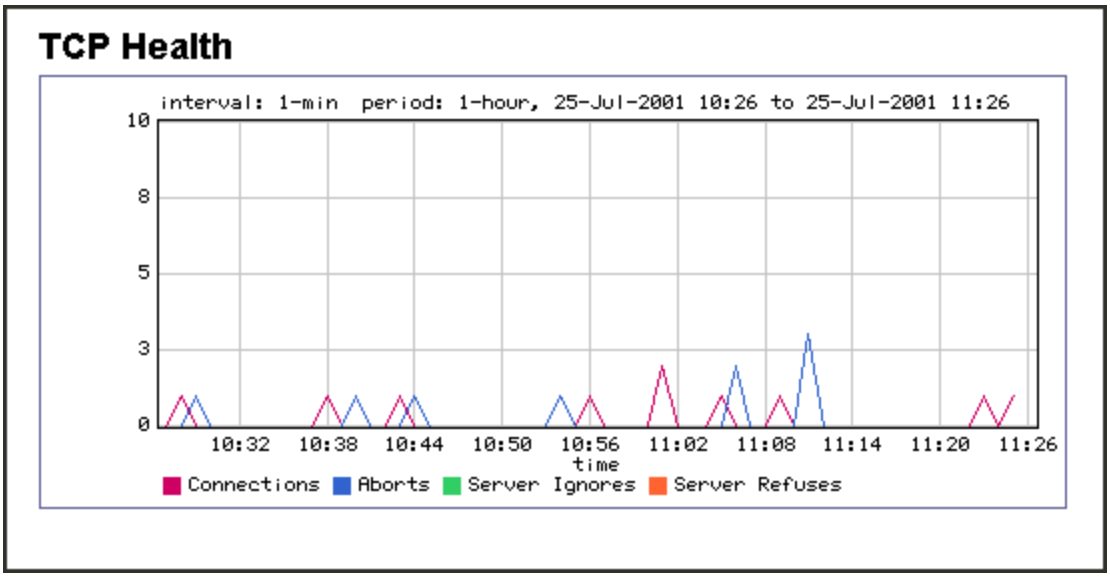

For example, suppose you created classes for each branch office and one of your offices complained about a server problem (such as slow speed, or many disconnects, or hung systems). By comparing TCP Health graphs of the various branch office classes, you can figure out what is causing the server problem at the one branch office.

For details on each of the measurement variables used in this graph (tcp-conn-inits, tcp-conn-aborts, tcp-conn-server-ignores, tcp-conn-server-refuses), see Measurement Variables.

Top 10 Children Classes Graph

The Top 10 Children Classes graph is a pie chart showing the relative portions of bandwidth allocated to the ten most active children classes of the selected class. This graph is similar to the Top 10 Classes graph except that this graph displays direct children only, while the Top 10 Class graph shows leaf classes (classes that don't have any children of their own). In other words, the Top 10 Children Classes graph allows you to graph a class' children, without grandchildren. (See a comparison of Top 10 Class and Top 10 Children Class graphs.)

This graph displays each child class' average bandwidth usage in bits per second and its percentage of the parent class' total bandwidth usage.

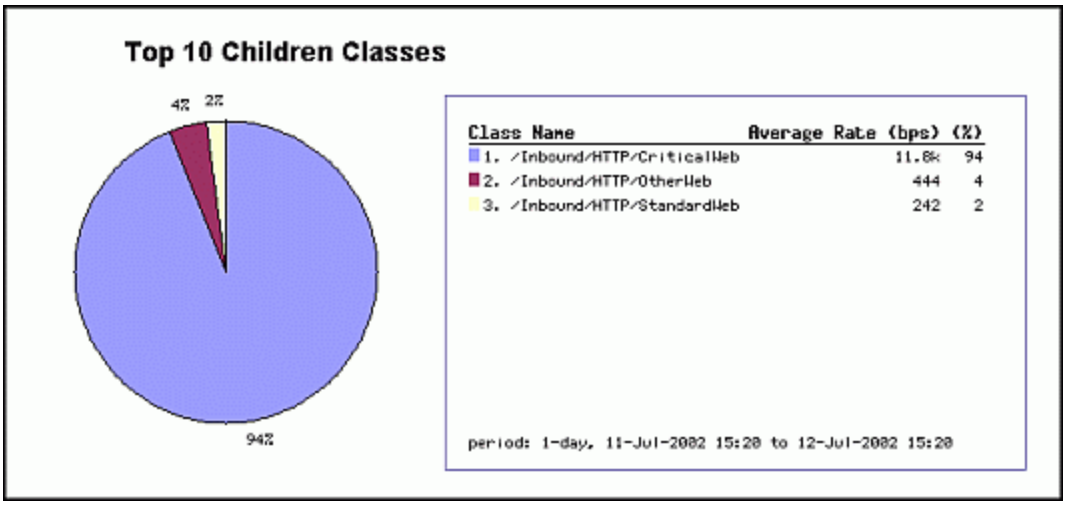

Top 10 Classes Graph

The Top 10 Classes graph is a pie chart showing the relative portions of bandwidth allocated to the ten most active classes associated with the selected class and all its descendents. This graph displays each class' average bandwidth usage in bits per second and its percentage of the total bandwidth used by the group.

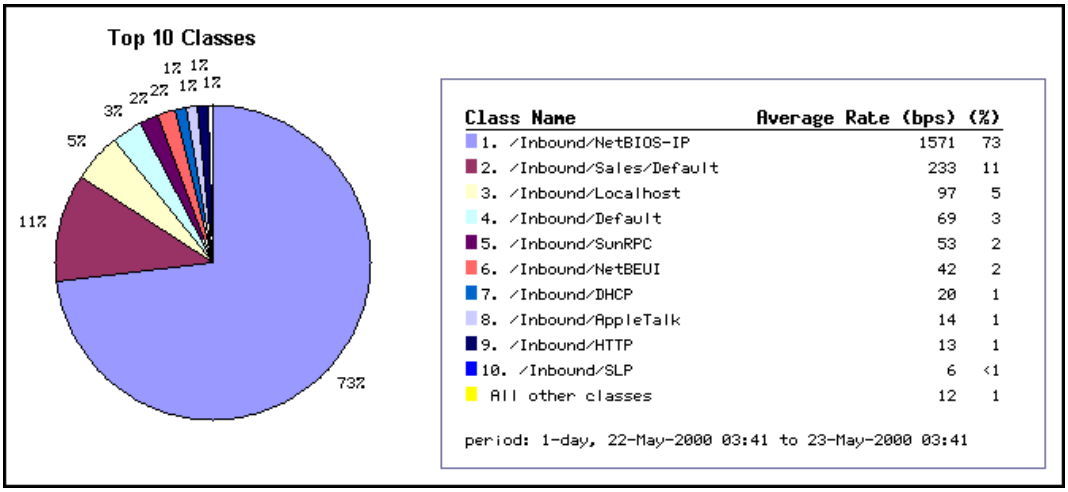

Class Name	Average Rate (bps)	(%)
1. /Inbound/NetBIOS-IP	1571	73
2. /Inbound/Sales/Default	233	11
3. /Inbound/Localhost	97	5
4. /Inbound/Default	69	3
5. /Inbound/SunRPC	53	2
6. /Inbound/NetBEUI	42	2
7. /Inbound/DHCP	20	1
8. /Inbound/AppleTalk	14	1
9. /Inbound/HTTP	13	1
10. /Inbound/SLP	6	<1
All other classes	12	1

period: 1-day, 22-May-2000 03:41 to 23-May-2000 03:41

Top 10 Children Classes vs. Top 10 Classes Graphs

These two pie charts graph the same time period for the same class (Inbound/HTTP). As you can see, Top 10 Children Classes graphs children only, while Top 10 Classes graphs children and grandchildren.

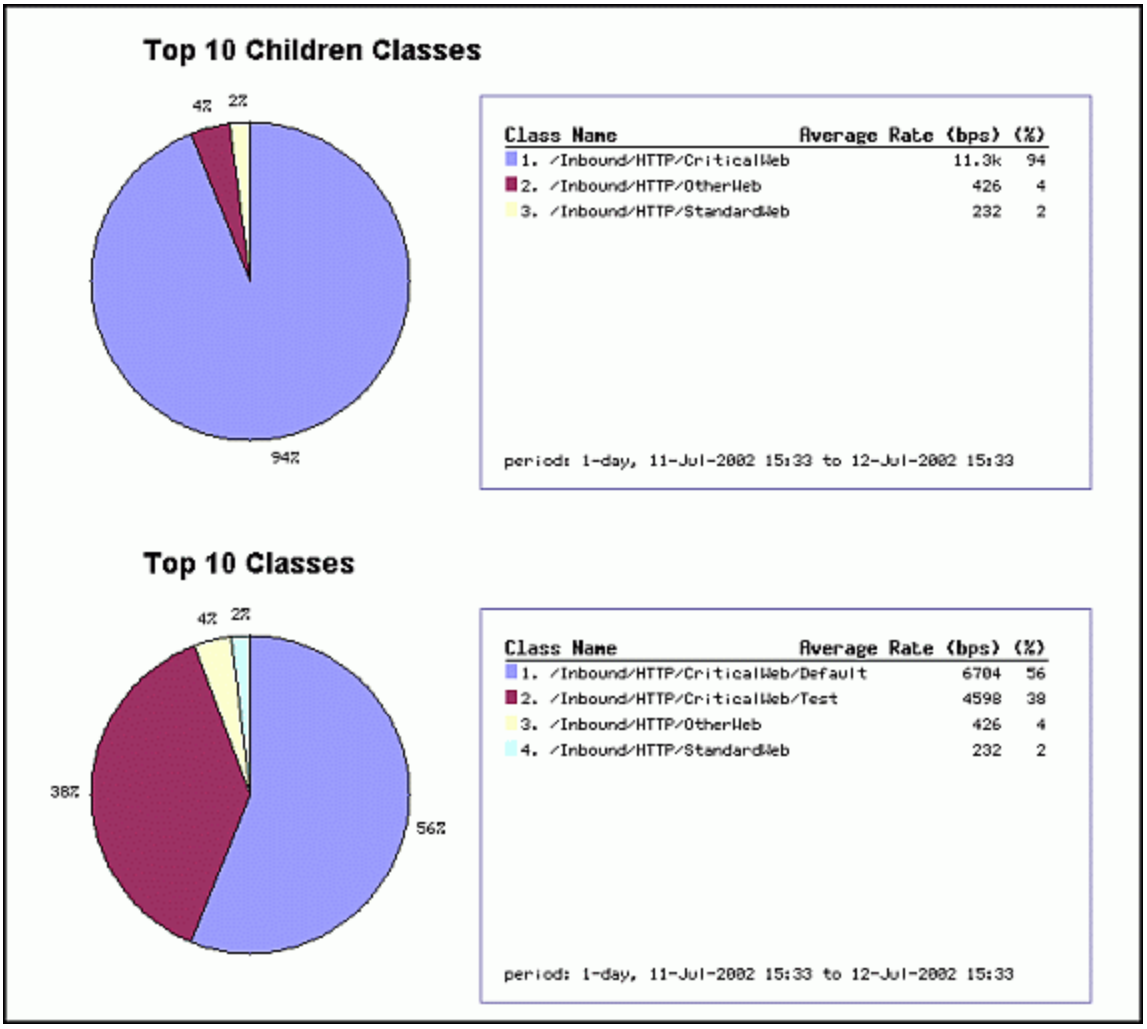

Top-10 Partitions Graph

The Top-10 Partitions graph is a pie chart showing the relative portions of bandwidth allocated to the ten most active partitions associated with the selected class or its descendents. This graph displays each partition's average bandwidth usage in bits per second and its percentage of the total bandwidth used by the group.

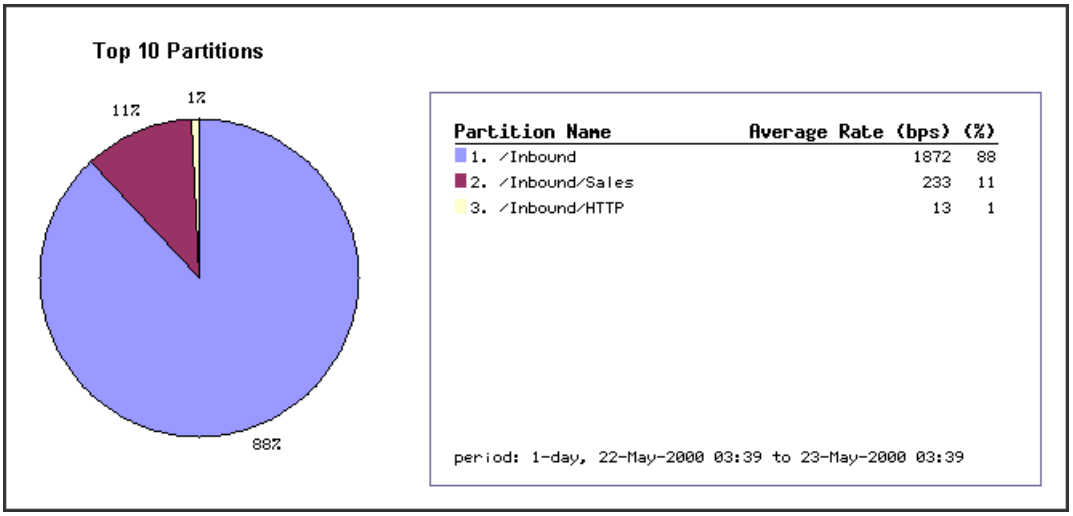

Note: Occasionally, the browser may display the colors incorrectly, causing discrepancies between the pie chart and the legend. To correct the colors, refresh the display by closing and reopening the window.

Transaction Delay Distribution Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

PacketShaper maintains histograms for network, server, and total delays for each class. The range of response times (0 to 25000+ milliseconds) is divided into 14 unequal transaction-time buckets. Each slot contains the number of transactions that are delayed with the specified number of seconds. For example, in the sample chart, about 140 transactions were delayed .1 seconds, 125 transactions were delayed .25 seconds, and 35 transactions were delayed .5 seconds. As each delay is calculated, PacketShaper updates its corresponding average and increases the count for the appropriate time slot.

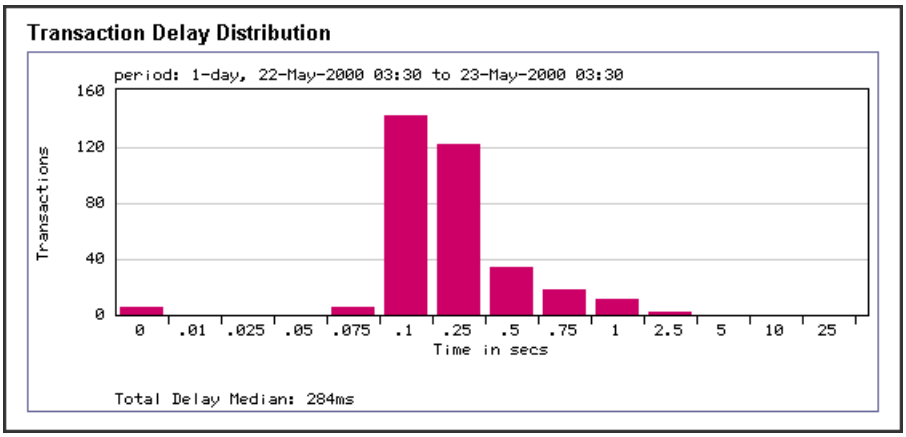

The Transaction Delay Distribution graph displays the number of transactions whose total delay falls into each of the 14 transaction-time buckets. The median delay is indicated on the bottom of the graph.

The transaction-time buckets and their respective counts provide efficient median estimates. A good median approximation is possible by using each time slot's count and making the assumption that the times within each group are distributed evenly.

Medians are presented to give a better idea of what most users are actually experiencing. Half the transactions' delays are below the median and half are above. While the average response time can be distorted by one very slow transaction, the median is not. For example, the average of this series (1, 2, 3, 4, 190) is 40, while the median is 3.

Transaction Delay Graph

> **Note:** This graph type is available only if RTM is available for the selected object.

The Transaction Delay graph shows a timeline of a traffic class' average response times in milliseconds. Three colored lines track the total, network, and server delays, so you can determine the source of any significant delays. The horizontal line represents the Total Delay Threshold as a reference point to determine the quality of the response times.

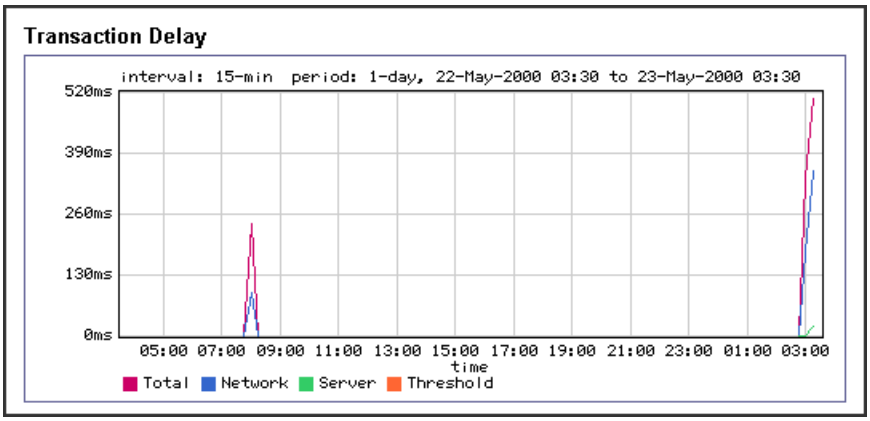

Allow URLs within a Blocked Category

Instructions to allow certain URLs for a category that is being blocked

There may be a category that you want to block on your network, but there could be one or more URLs within that category that you want to allow users access to.

Steps:

1. Block the traffic in a category. We'll refer to this class as Blocked_Category in this example.
2. Create a host list (in Sky or Advanced UI) that identifies the URLs you want to allow from that category.
3. With the Blocked_Category class selected, create a class (in Sky or Advanced UI) called Allow_List with a matching rule for the host list created in step 2.

   > **Note:** The Allow_List class should be a child class of the Blocked_Category class.

4. Assign a rate (in Sky or Advanced UI) or priority (in Sky or Advanced UI) policy to the Allow_List class.
5. (Advanced UI only) Repeat the above steps for the Outbound direction.

   > **Note:** the Sky UI has the option to create classes in both directions, eliminating the need to repeat the steps.

6. Make sure shaping is enabled.

The class tree should look similar to the following:

```
Blocked_Category
 Allow_List <--Priority(3) policy
      Default <-- Never-Admit policy
```

The never-admit policy was originally assigned to the Blocked_Category class, but when you created the Allow_List child class, the policy automatically shifted to the Blocked_Category/Default class. Thus, the URLs in the Allow_List get assigned the Priority(3) policy and are allowed; all other URLs in the category get assigned the never-admit policy and are blocked.

Allow Users Access to a Blocked Category

Instructions to allow certain users access to a category that is being blocked

There may be a category that you generally want to block on your network, but there could be one or more users that you want to allow access to that category. For example, you may want to open up access to an administrator's static IP address or a subnet of users.

Steps:

1. Block the traffic in a category. We'll refer to this class as Blocked_Category in this example.

2. Create a host list (in Sky or Advanced UI) that identifies the static IP addresses and/or subnets of the users who you want to allow access to the blocked category.

3. With the Blocked_Category class selected, create a class (in Sky or Advanced UI) called Allow_List with an inside matching rule for the host list created in step 2.

   > **Note:** The Allow_List class should be a child class of the Blocked_Category class.

4. Assign a rate (in Sky or Advanced UI) or priority (in Sky or Advanced UI) policy to the Allow_List class.

5. (Advanced UI only) Repeat the above steps for the Outbound direction.

   > **Note:** the Sky UI has the option to create classes in both directions, eliminating the need to repeat the steps.

6. Make sure shaping is enabled.

The class tree should look similar to the following:

```
Blocked_Category
 Allow_List <--Priority(3) policy
 Default <-- Never-Admit policy
```

The never-admit policy was originally assigned to the Blocked_Category class, but when you created the Allow_List child class, the policy automatically shifted to the Blocked_Category/Default class. Thus, the URLs in the Allow_List get assigned the Priority(3) policy and are allowed; all other URLs in the category get assigned the Never-Admit policy and are blocked.

Block URLs within an Allowed Category

Instructions to block certain URLs within a category that is allowed

There may be a category that you want to generally allow on your network, but there could be some exceptions–you want to block users' access to one or more URLs within that category.

Steps:

1. Make sure the traffic tree contains a class for the category you want to allow, and that this class doesn't have a blocking (never-admit or discard) policy. We'll refer to this class as Allowed_Category in this example.

2. Create a host list (in Sky or Advanced UI) that identifies the URLs you want to block.

3. With the Allowed_Category class selected, create a class (in Sky or Advanced UI) called Deny_List with a matching rule for the host list created in step 2.

   > **Note:** The Deny_List class should be a child class of the Allowed_Category class.

4. Assign a never-admit policy (in Sky or Advanced UI) to the Deny_List class.

5. (Advanced UI only) Repeat the above steps for the Outbound direction.

   > **Note:** the Sky UI has the option to create classes in both directions, eliminating the need to repeat the steps.

6. Make sure shaping is enabled.

The class tree should look similar to the following:

```
Allowed_Category
 Deny_List <--Never-Admit policy
 Default <-- no policy (uses default policy of Priority(3))
```

The URLs in the Deny_List get assigned the never-admit policy and are blocked; all other URLs in the category get assigned the default policy and are allowed.

Block Users Access to an Allowed Category

Instructions to block certain users access to a category that is allowed

There may be a category that you want to generally allow on your network, but there could be some exceptions—you want to block certain users' access to that category. The users are defined by IP address or subnet.

Steps:

1. Make sure the traffic tree contains a class for the category you want to allow, and that this class doesn't have a blocking (never-admit or discard) policy. We'll refer to this class as Allowed_Category in this example.

2. Create a host list (in Sky or Advanced UI) that identifies the static IP addresses and/or subnets you want to block.

3. With the Allowed_Category class selected, create a class (in Sky or Advanced UI) called Deny_List with an inside matching rule for the host list created in step 2.

   > **Note:** The Deny_List class should be a child class of the Allowed_Category class.

4. Assign a never-admit policy (in Sky or Advanced UI) to the Deny_List class.

5. (Advanced UI only) Repeat the above steps for the Outbound direction.

   > **Note:** the Sky UI has the option to create classes in both directions, eliminating the need to repeat the steps.

6. Make sure shaping is enabled.

The class tree should look similar to the following:

```
Allowed_Category
  Deny_List <--Never-Admit policy
  Default <-- no policy (uses default policy of Priority(3))
```

The IP addresses in the Deny_List get assigned the never-admit policy and are blocked; all other IP addresses in the category get assigned the default policy and are allowed.

Remove Units from PolicyCenter

PacketShapers can operate in either **shared mode** or **local mode**. A unit in local mode uses a configuration derived only from the files on the unit itself, while units in shared mode use PolicyCenter for configuration information. *To remove a unit from PolicyCenter, you must first switch the unit to local mode.*

While the unit is in shared mode, PolicyCenter continually updates the configuration files on the unit's system disk to match its PolicyCenter configuration. Therefore, after you switch from shared mode back to local mode, the unit's local configuration will be the same as its previous PolicyCenter configuration. Be aware that simply switching a unit to local mode does not remove it completely from PolicyCenter– you must also delete the unit configuration from PolicyCenter.

Switch the Unit to Local Mode

To switch a unit to local mode:

1. Access the unit's browser interface, using a browser and the IP address of the unit.
2. Go to the Advanced UI and click the **Setup** tab.
3. From the **Choose Setup Page** list, select **PolicyCenter access**.
4. Select **local mode**.

   A pop-up window will ask you to confirm that you want to switch the unit to local mode.
5. Click **OK**.

Delete the Unit Configuration from PolicyCenter

1. In PolicyCenter, click the **Configurations** tab.
2. Make sure the **Show unit configs** check box is selected.
3. Select the unit configuration that you switched to local mode.
4. Click the **Operations** tab.
5. Under Delete Configuration, click the **Delete** button.

Print and Update an Event Graph

The event graph (below) displays three colored lines that represent the following event statistics:

- Blue line – The recorded statistics for the event's measurement variable
- Red line – The threshold value for the event. The event is triggered when the threshold condition is met
- Green line – The re-arm value for the event. The event is re-armed after a threshold condition is met, so that another event can be triggered

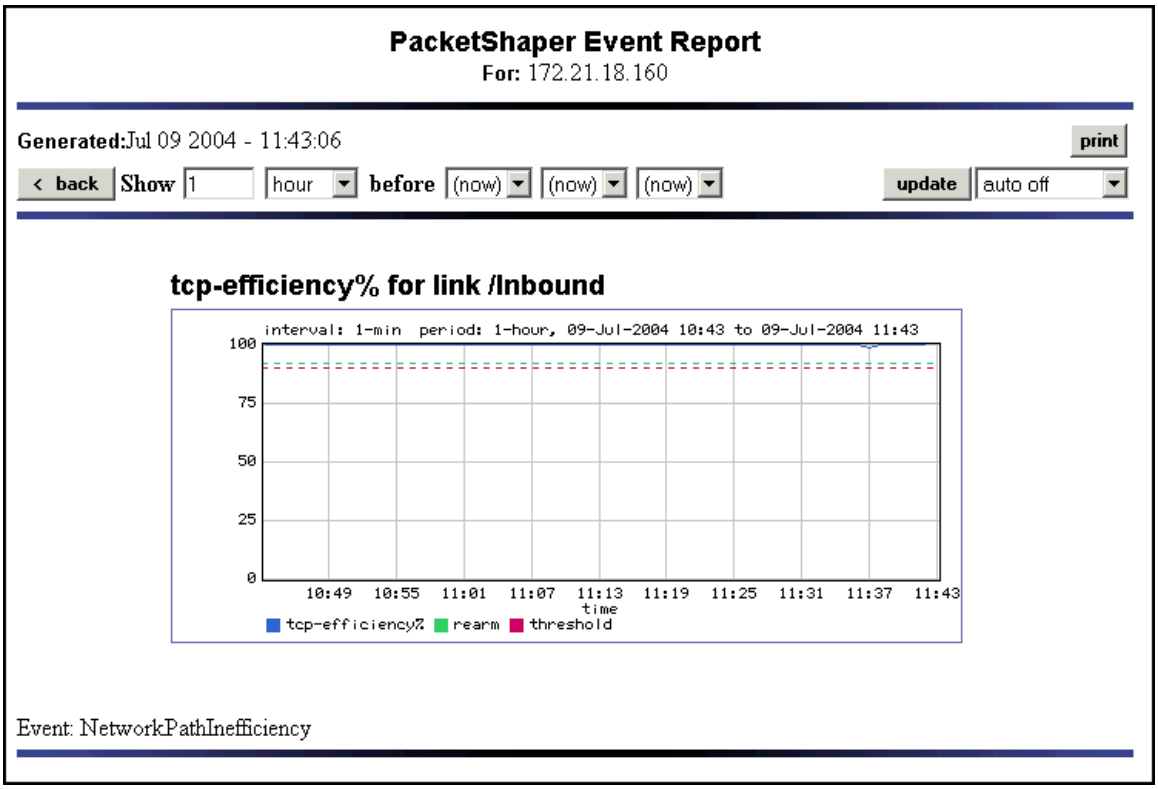

By default, the event graph analyzes traffic over the last hour, but the time period and end date are user-customizable. The ability to create reports for a specific time period is useful for creating comparison reports. For example, you can print one report that analyzes traffic over the last week (a period of one week, with **now** as the end date and time), and then print a second report that analyzes traffic for the week before that (a period of one week, with a specific date specified for the end date the last day you want to analyze).

To specify a different time period:

1. In the **Show** field, specify the number of minutes, hours, days, weeks, or months for which you want traffic analyzed.

For example, by entering **2** in the **Show** field and selecting **week** for the unit of time, you create a report with the data PacketShaper collected over the last two weeks. The minimum period you can specify is 1 minute; the maximum is 60 days (or 2 months).

2. To choose a specific end date for the report, select the month and day from the **Before** drop-down lists. If you want the current date as the end date, leave the default values of **now**.

> **Note:** Although you can select a date from the previous 12 months, the unit may not store data that far back– the amount of historical data that is stored depends on how many traffic classes you have in your tree. The unit stores measurement data for a minimum of 31 days in hourly intervals and at least 1 day of data in one-minute intervals.

3. To choose a specific end time for the report, select the time from the drop-down list. If you want the current time as the end time, leave the default values of **now**.

4. Click **update** to update the statistics for the new time period.

> **Note:** To update the graph at regular intervals, select an interval from the drop-down list located next to the **update** button.

5. To print the report, click **print**.

Set Up Syslog

The logging feature gives administrators a way to centrally log and analyze configuration events and system error messages. For example, if you are using RADIUS authentication, each failed login attempt will be sent to the defined Syslog server.

Syslog can be used in conjunction with the adaptive response and user events feature. For example, when you register an event, you will be asked if you want to send events to Syslog; you can define and register an event that sends a message to a Syslog server when retransmissions rise to 30 percent of your network activity.

To set up logging, you need to:

1. Define one or more Syslog servers.

2. Enable the logging feature.

Define the Syslog Server

To define a Syslog server, use the **setup syslog add** command from the command-line interface as described below. You can add up to four Syslog servers.

To define a Syslog server:

1. Access the CLI.

2. At the command-line prompt, use the following command:

```
setup syslog add host:<ipaddress> [output:<facility>,<level>] [port:<portnum>] [datetime]
```

host: *<ipaddress>*	The Syslog server IP address for example, *host:10.7.38.100*
output: *<facility>* ,*<level>*	The facility and severity level for example, *output:local1,6* Up to three outputs can be specified. The default facility is local4 and the default level is 7. PacketShaper user events are at severity level 6; if you want to capture them with Syslog, you must set the level to 6 or 7. See Facility Types and Severity Levels below for lists of the valid facility types and levels.
port: *<portnum>*	The port number of the Syslog server; if the port isn t specified, port 514 is used
datetime	Include the date and time in the message; the date and time are not included unless you specify the **datetime** parameter

For example:

```
setup syslog add host:10.7.38.100 output:local4,3 datetime
```

Facility Types

You can enter the keyword or value specified in the following table.

Description	Keyword	Value
Kernel	kern	0
User Processes	user	1
Electronic Mail	mail	2
Background System Processes	daemon	3
Authorization	auth	4
System Logging	sysl	5
Printing	lpr	6
Usenet News	news	7
Unix-to-Unix Copy Program	uucp	8
Clock Daemon	clkd	9
Security	sec2	10
FTP Daemon	ftpd	11
NTP Subsystem	ntp	12
Log Audit	audi	13
Log Alert	alert	14
Clock Daemon	clkd2	15
For Local Use	local0-local7	16-23

Severity Levels

You can enter the keyword or value specified in the following table. Set the level to specify which messages to suppress to the Syslog server. For example, setting the severity level to 3 allows messages with levels 0 - 3 and suppresses messages with levels 4 - 7. If you don't specify a severity level, 7 is used. With the default severity level, messages of all levels will get sent to the Syslog server.

Description	Keyword	Value
System unusable	emerg	0
Take immediate action	alert	1
Critical condition	crit	2
Error message	err	3
Warning message	warn	4
Normal but significant condition	notice	5
Informational (includes PacketShaper user events)	info	6
Debug message	debug	7

At the "warn" level, PacketShaper will send the following types of messages to the Syslog server:

- Login failed
- Hard drive status
- Measurement Engine status

User events that are configured to send a syslog message when a threshold is crossed are sent at the info severity level (6). See event register for more information on configuring an event to send a syslog message.

Enable the Logging Feature

To enable the logging feature so that messages will be sent to the defined Syslog server(s), type the following command at the command-line prompt:

```
setup syslog state on
```

Update the Event Statistics

The statistics in the *Monitor Events* window are captured the moment you choose **Events** on the **Monitor** drop-down list. To get a new snapshot, click **update** at the top of the screen.

You can set the *Monitor Events* window to automatically update at regular intervals. To enable auto-update:

1. Click **Auto**. You are prompted for an interval.

2. Enter the desired number of seconds between updates, and click **OK**.

To stop automatic updating, click **Stop Auto**.

To increase the update interval, do one of the following:

- Click **Auto**. Each time you click **Auto**, three seconds are added to the interval.
- Click **Stop Auto** to disable automatic updates, and then click **Auto** to enter a new interval.

Disable Access-Link Monitoring

When you turn off high availability, all your configuration settings are retained, but PacketShaper will not monitor the status of the interfaces and routers, nor will it adjust the access link size and resize partitions if a link goes down.

To disable access-link monitoring:

1. Click the **Setup** tab.

2. From the **Choose Setup Page** list, select **High Availability**. The *High Availability* screen is displayed.

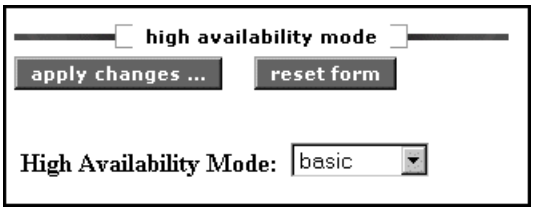

3. In the **High Availability Mode** field, choose **off**.

4. Click **apply changes**.

Define Host Accounting Categories

With the host accounting feature, you can set up categories and then assign traffic classes to the appropriate category. Categories give you the ability to add another dimension to host data. In addition to tracking how many bytes have been sent to and received from an IP address or subnet, you can see a breakdown of the type of traffic associated with each user. For example, suppose you want to bill a customer for its web traffic. You can set the groundwork for this by creating a category called "web" and assigning the Inbound/HTTP and Outbound/HTTP classes to it.

> **Tip:** Give careful up-front thought to your category names. You must create all your categories at once, and if you later want to modify your category list, you must redefine all the categories.

1. Access the command-line interface.
2. Use the host accounting categories command to create the categories.
3. Use the class category command to assign traffic classes to each category.
4. Enable host accounting.

Enable Host Accounting

When you enable host accounting, you define the host location, the number of minutes between each recorded sample, and the maximum number of samples that can be stored in the host database.

1. Access the command-line interface.

2. Use the host accounting enable command to enable host accounting.

After you enable host accounting host accounting measurement data will not begin recording again until the next full interval. For example, assume *<interval-minutes>* is 2 and you reset the unit. When you give the **measure show** command, the message indicates "Measurement engine is waiting until 15:03 to start." At 15:03 all other measurement groups will begin recording, but host accounting will not begin recording until 15:04 (the next interval).

To see how many samples have been recorded, use the measure show command.

To view the recorded data, see Retrieve Host Accounting Data.

To check your host accounting settings, use the host accounting show command.

Retrieve Host Accounting Data

After host data has collected for a sufficient amount of time, you can view the data on the screen and/or save it into a file that can be imported into another program for additional calculations (such as totaling and averaging) and billing. You have a lot of flexibility when retrieving host data. You can display host data for all hosts, for a certain IP address, a subnet, or a host list. In addition, you specify the period for which you want to retrieve data a range of dates or relative dates (for example, -7 for 7 days ago).

> **Note:** Before you can retrieve host accounting data, you need to enable the feature.

1. Access the command-line interface.

2. Use the host accounting retrieve command to display the host data. To save the data in a comma-separated value (.CSV) file, use the **into** ***<file>*** parameter.

3. Use the File Browser to download the .CSV file to a local workstation.

4. Import the .CSV file into a third-party application (such as Microsoft Excel).

Create Custom Reports

The Statistical Graphing custom report feature allows you to generate reports that graph any of the measurement variables supported for the selected class. For a description of each measurement variable, see Measurement Variables.

To create a custom report:

1. Click the **Report** tab.

2. In the Tools section on the left, click **Statistical Graphing**.

> **Note:** In order to use the Statistical Graphing Tool, you must have the Adobe Flash Player plug-in). If you do not have it installed, you will be prompted to download and install it before you can proceed.

The Historical Overview window is displayed. **show screen...**

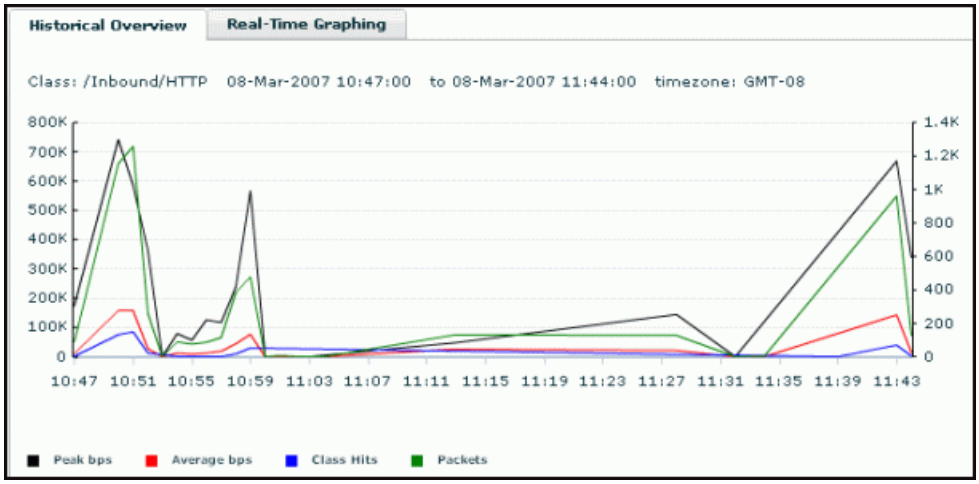

3. In the Class Tree, select the class for which you want to create a report. As soon as you click on a class name, the corresponding graph is displayed in the right window pane.

   Click the ▶ arrow next to the Inbound or Outbound folder to expand the tree. If you have many classes to navigate through, you can filter the class tree by entering a search string in the search field. Only classes containing the string will be displayed.

4. Define the time period for the report by selecting **Week**, **Day**, or **Hour**.

   By default, data from the last hour is graphed. When you change the time period, the new graph renders instantly.

5. Click the **Reports** tab. **show screen...**

6. Define the time window for the report:

- Slide the purple bar below the graph to the left or right to shift the report window earlier or later. **show screen...**

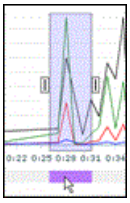

- Increase or decrease the window size by dragging the window handles. The report you generate will cover the period contained within the purple window. **show screen...**

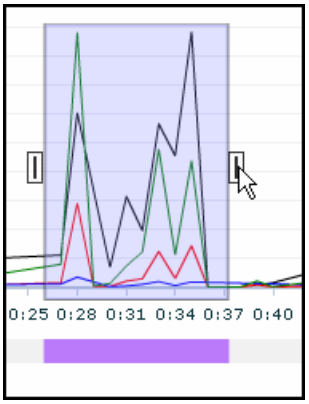

- Shift the reporting time period by increments of one (week, day, or hour depending on the time period you selected on the **Overview Control Panel** tab) by clicking the left- or right-arrow below the graph.

7. Select the measurement variables you want to graph on the custom report from the **Metric Type** list. Up to 56 variables can be selected. You can also filter the list of variables by entering a search string in the search field. Only measurement variables containing the string will be displayed.

- Ctrl+Click to select a set of non-contiguous measurement variables.
- Shift+Click to select a set of contiguous measurement variables.

For a description of each variable, see Measurement Variables.

8. To label the report, enter a title in the **Report Title** field.

9. To generate the report, click **Create Custom Report**. The report opens in a new browser window.

10. To print the report, click **File** and select a print option:

- To print the report on a single page, select **Print to One Page**.
- To print the report on multiple pages, select **Print**.

11. To save the report in a file format that can be imported into another application, click **File** in the report window and then select one of the following file formats:

- Select **as CSV** to save the report as a comma-separated values file, which can be opened with a spreadsheet application, such as Microsoft Excel, or imported into a database application.
- Select **as TSV** to save the report as a tab-separated values file, which can be opened with a spreadsheet application, such as Microsoft Excel, or imported into a database application.
- Select **as XML** to save the report as an Extensible Markup Language file, which can be imported directly into Oracle 8i.

Additional Resources

For information on analyzing data shown in the Statistical Graphing Tool, see the following PacketGuide Recommendations:

Analyze an Application's Behavior
Identify Performance Saboteurs
Access Wasted Network Capacity

Create Standard Reports

The Statistical Graphing standard report feature allows you to generate reports containing one or more of the following preconfigured graphs for analyzing bandwidth usage, efficiency, and TCP connections for a selected class:

- Class Utilization with Peaks
- Bytes Transmitted
- Packets Transmitted
- TCP Health
- Connection Retransmissions

For descriptions and examples of these and other graphs, see Preconfigured Graphs.

To create a standard report:

1. Click the **Report** tab.

2. In the Tools section on the left, click **Statistical Graphing**.

> **Note:** In order to use the Statistical Graphing Tool, you must have the Adobe Flash Player plug-in. If you do not have it installed, you will be prompted to download and install it before you can proceed.

The Historical Overview window is displayed. **show screen...**

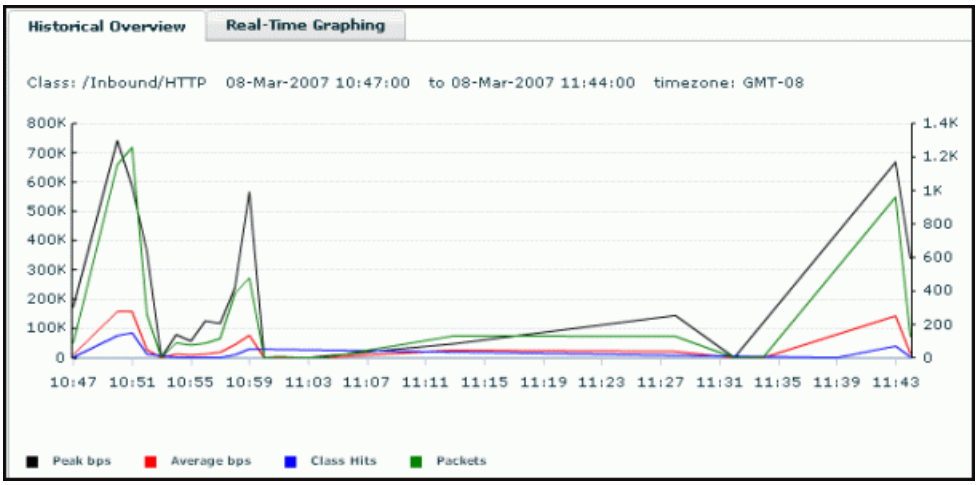

3. In the Class Tree, select the class for which you want to create a report. As soon as you click on a class name, the corresponding graph is displayed in the right window pane.

Click the ▶ arrow next to the Inbound or Outbound folder to expand the tree. If you have many classes to navigate through, you can filter the class tree by entering a search string in the search field. Only classes containing the string will be displayed.

4. Define the time period for the report by selecting **Week**, **Day**, or **Hour**.

By default, data from the last hour is graphed. When you change the time period, the new graph renders instantly.

5. Click the **Reports** tab. **show screen...**

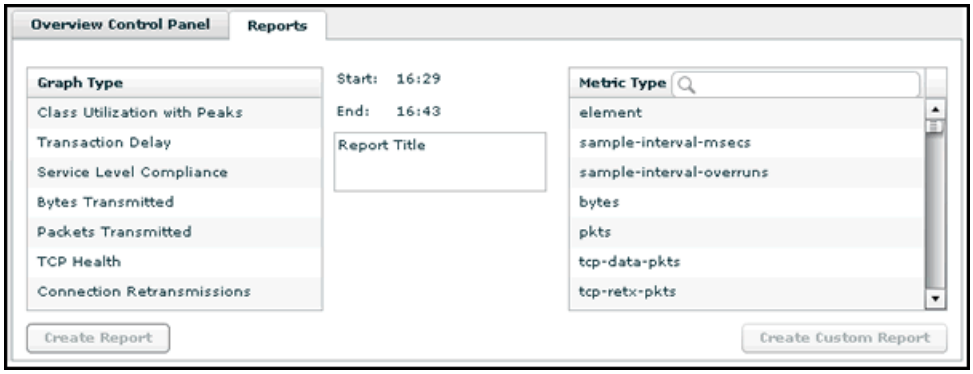

6. Define the time window for the report:

- Slide the purple bar below the graph to the left or right to shift the report window earlier or later. **show screen...**

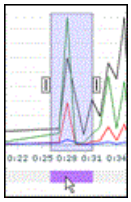

- Increase or decrease the window size by dragging the window handles. The report you generate will cover the period contained within the purple window. **show screen...**

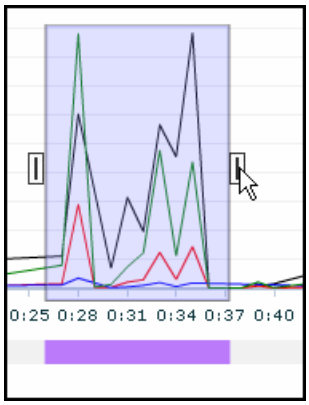

- Shift the reporting time period by increments of one (week, day, or hour depending on the time period you selected on the **Overview Control Panel** tab) by clicking the left- or right-arrow below the graph.

7. Select the graphs to include on the report by clicking the graph name in the Standard Graphs list on the **Reports** tab. You can select up to seven graphs to include in the report:

- Ctrl+Click to select a set of non-contiguous graphs.
- Shift+Click to select a set of contiguous graphs.

8. To label the report, enter a title in the **Report Title** field.

9. To generate the report, click **Create Report**. The report opens in a new browser window.

10. To print the report, click **File** and select a print option:

- To print the report on a single page, select **Print to One Page**.
- To print the report on multiple pages, select **Print**.

11. To save the report in a file format that can be imported into another application, click **File** in the report window and then select one of the following file formats:

- Select **as CSV** to save the report as a comma-separated values file, which can be opened with a spreadsheet application, such as Microsoft Excel, or imported into a database application.
- Select **as TSV** to save the report as a tab-separated values file, which can be opened with a spreadsheet application, such as Microsoft Excel, or imported into a database application.
- Select **as XML** to save the report as an Extensible Markup Language file, which can be imported directly into Oracle 8i.

Additional Resources

For information on analyzing data shown in the Statistical Graphing Tool, see the following PacketGuide Recommendations:

Analyze an Application's Behavior
Identify Performance Saboteurs
Access Wasted Network Capacity

Delete Reports

Custom reports are listed in the My Reports section of the **Report** tab. **show screen...**

To delete a custom report:

1. Click the **Report** tab.

2. Click **delete report** in the left window pane. The *Delete Reports* window appears. **show screen...**

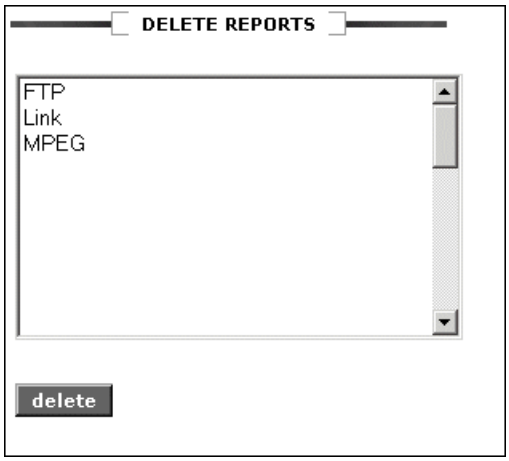

3. Select the report you want to remove.

4. Click **delete**.

IPv6 Traffic Summary Report

View the IPv6 Traffic Summary to see how much IPv6 traffic is on your network and compare it to other types of traffic—IPv4 and non-IP. This report contains two line graphs (Inbound and Outbound) that show the utilization of various types of IP traffic during the specified time period. In addition, the report includes a summary table of statistics, allowing you to see the breakdown of IPv4, IPv6, non-IP, and total bytes sent and received during the time period.

> **Note:** It is not necessary to have IPv6-only classes in your traffic tree for this graph to populate; this data is tracked as a separate variable for all classes at the link level.

To display the IPv6 Summary:

1. Click the **Report** tab.

2. Click **IPv6 Summary**. The IPv6 Traffic Summary report displays in the current window.

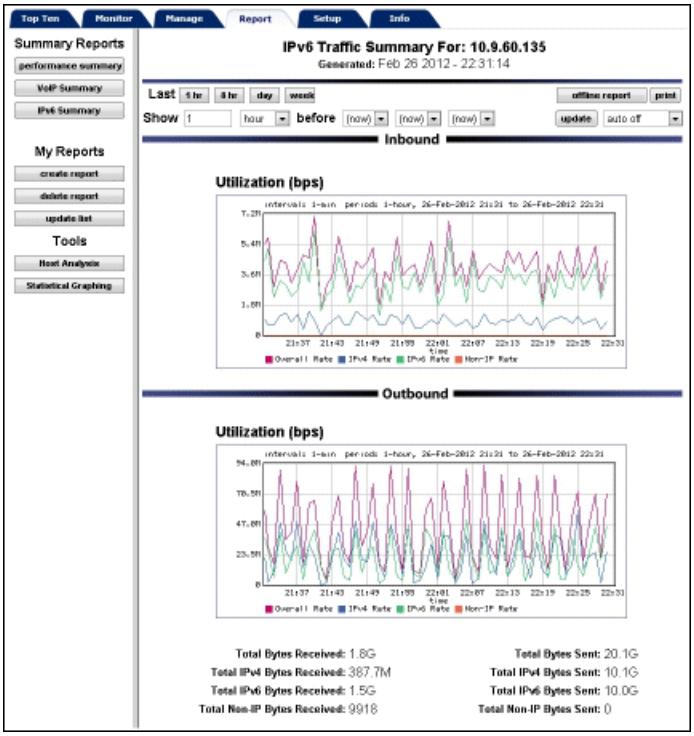

3. Adjust the time period and end date, if desired.

4. The default report contains data from the current PacketShaper. If you have enabled Standby, you can view the IPv6 Traffic Summary of the other partners in the high availability cluster. Choose the partner's IP address from the **Standby Partner** list.

5. If the report has been displayed for awhile, click **update** to retrieve the most current data.

6. To print the report, click **print**.

The following statistics are graphed in the Inbound and Outbound Utilization line graphs:

Statistic	Description
Overall Rate	Rate of all traffic, averaged over the interval, in bits per second. This is an exponentially weighted moving average of five seconds of data throughput.
IPv4 Rate	Rate of IPv4-only traffic, averaged over the interval, in bits per second.
IPv6 Rate	Rate of IPv6-only traffic, averaged over the interval, in bits per second.
Non-IP Rate	Rate of non-IP traffic, averaged over the interval, in bits per second.

Select the Time Period

By default, the reports in the Top Ten window analyze network traffic over the last hour, but the time period and end date are user-customizable. The ability to create reports for a specific time period is useful for creating comparison reports. For example, you can print one report that analyzes traffic over the last week (a period of one week, with **now** as the end date and time), and then print a second report that analyzes traffic for the week before that (a period of one week, with a specific date specified for the end date the last day you want to analyze). **show screen...**

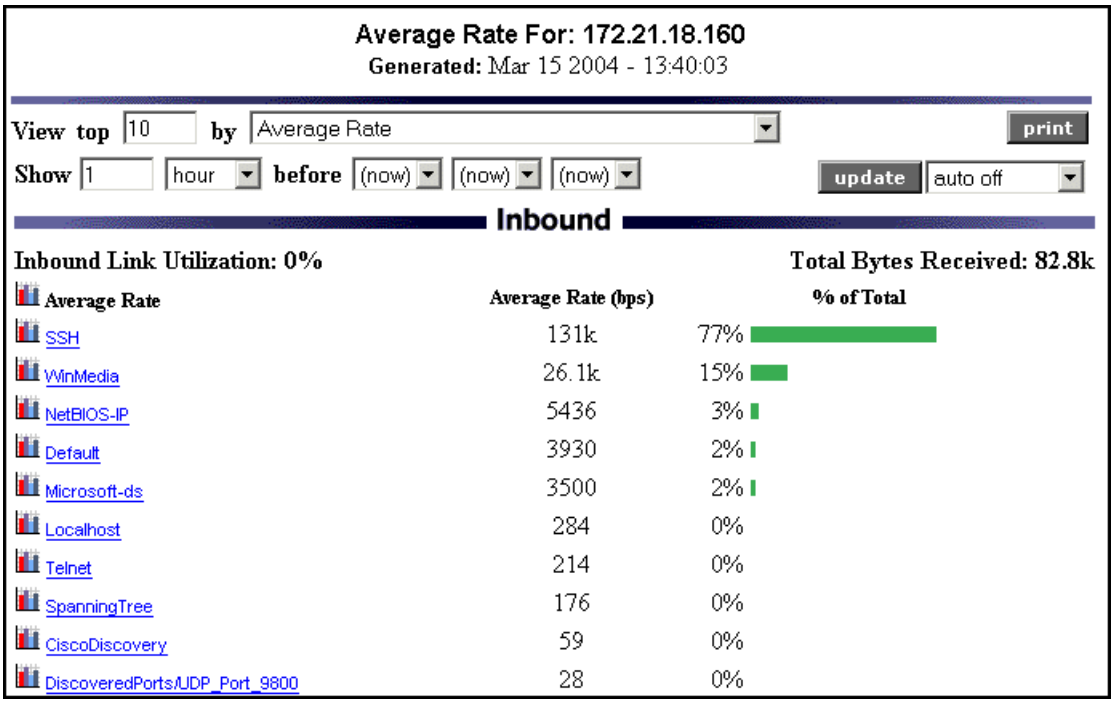

To specify a different time period:

1. In the **Show** field in the Top Ten window, specify the number of minutes, hours, days, weeks, or months for which you want traffic analyzed.

   For example, by entering **2** in the **Show** field and selecting **week** for the unit of time, you create a report with the data PacketShaper collected over the last two weeks. The minimum period you can specify is 1 minute; the maximum is 60 days (or 2 months).

2. To choose a specific end date for the report, select the month and day from the **Before** drop-down lists. If you want the current date as the end date, leave the default values of **now**.

> **Note:** Although you can select a date from the previous 12 months, the unit may not store data that far back the amount of historical data that is stored depends on how many traffic classes you have in your tree. The unit stores measurement data for a minimum of 31 days in hourly intervals and at least 1 day of data in one-minute intervals.

3. To choose a specific end time for the report, select the time from the drop-down list. If you want the current time as the end time, leave the default values of **now**.

4. Click **update** to update the statistics for the new time period.

> **Note:** To update the statistics at regular intervals, select an interval from the drop-down list located next to the **update** button.

5. To print the report, click **print**.

Select a Different View

In the Top Ten window, you can view a variety of reports, such as a report of network efficiency. You can also change the number of classes displayed. **show screen...**

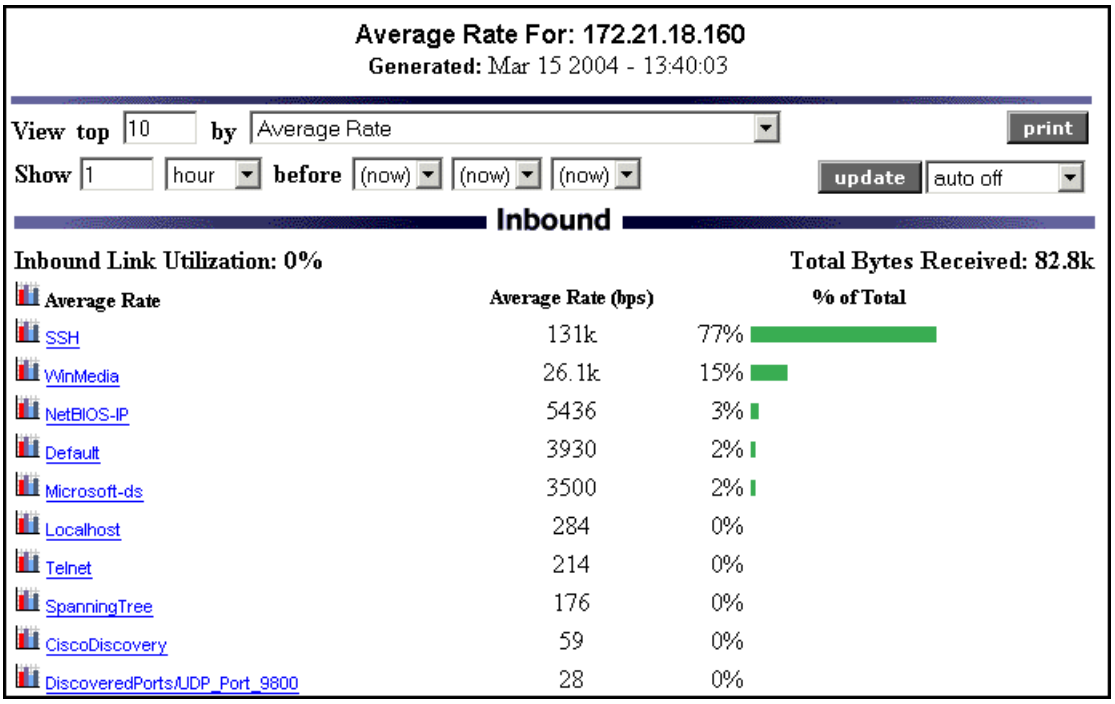

To select a different view in the *Top Ten* window:

1. Click the **Top Ten** tab.

2. In the **View top** field, enter the number of classes you want to view in the report (default is 10). To view all classes, clear this field so that it's blank.

> **Note:** If you don't want to change the view (the next step), you must click **update** to see the revised list.

3. Select the desired view:

View	Description
Average Rate	The classes with the highest average bandwidth rate in the designated time interval; measured in bits per second (the default report type)
Peak Rate	The classes with the highest peak bandwidth rate in the designated time interval; measured in bits per second
Total Bytes	The classes with the highest total bandwidth consumed in the designated time interval; measured in bytes
Network Efficiency	The classes with the most retransmissions (that is, the lowest network efficiency)

4. Make sure the time period is appropriate and change it if necessary. See Select the Time Period.

5. To print the new report, click **print**.

Getting More Details About Top Classes

The reports on the Top Ten page include a few ways to drill down and get more information about each class.

- Next to each class name is a graph icon . When you click this icon, a report containing two or three graphs pertaining to the class appears in its own window. For example, if you are viewing the Network Efficiency report and click the graph icon next to the HTTP class, a Class Details report appears with an Average and Peak Rates graph and a Network Efficiency graph of the HTTP class. Note that the graph icon appears to the left of the column heading (Top Ten - Average Rate, for instance) as well. Clicking this icon displays stacked charts of the top ten classes.

- Each class name is a hyperlink. When you click the class name, the *Traffic Class* page for that class appears so that you can check settings or make modifications.

Print and Update a Custom Report

1. To make sure the report is displaying the most current data:

- Click **update**.

or

- Choose one of the auto-updating options: **every 1 min**, **every 5 min**, **every 15 min**, or **every hour**.

2. If you want to view a report of the last eight hours, day, or week, use one of the preset buttons:

And then click **print**. (You can skip the remaining steps.)

3. To specify a different time period, enter the number of minutes, hours, days, weeks, or months in the **Show** field.

For example, by entering **2** in the **Show** field and selecting **week** for the unit of time, you create a report with the data PacketShaper collected over the last two weeks. The minimum period you can specify is 1 minute; the maximum is 60 days (or 2 months).

4. To choose a specific end date for the report, select the month and day from the **before** drop-down lists. If you want the current date as the end date, leave the default values of **now**.

> **Note:** Although you can select a date from the previous 12 months, PacketShaper may not store data that far back the amount of historical data that is stored depends on how many traffic classes you have in your tree. The unit stores measurement data for a minimum of 31 days in hourly intervals and at least 1 day of data in one-minute intervals.

5. To choose a specific end time for the report, select the time from the drop-down list. If you want the current time as the end time, leave the default values of **now**.

6. To print the report, click **print**.

Retrieve Saved Reports

Saved custom reports are listed as buttons in the left pane of the report window. **show screen...**

To display a saved report:

1. Click the **Report** tab.

2. In the My Reports list in the left window pane, select the name of the report you want to view. The report appears in the right window pane.

3. *Optional:* In the **Show** field, specify the number of minutes, hours, days, weeks, or months for which you want traffic analyzed.

   For example, by entering**2** in the **Show** field and selecting **week** for the unit of time, you create a report with the data PacketShaper collected over the last two weeks. The minimum period you can specify is 1 minute; the maximum is 60 days (or 2 months).

4. *Optional:* To choose a specific end date for the report, select the month and day from the **Before** drop-down lists. If you want the current date as the end date, leave the default values of **now**.

> **Note:** Although you can select a date from the previous 12 months, the unit may not store data that far back the amount of historical data that is stored depends on how many traffic classes you have in your tree. The unit stores measurement data for a minimum of 31 days in hourly intervals and at least 1 day of data in one-minute intervals.

5. *Optional:* To choose a specific end time for the report, select the time from the drop-down list. If you want the current time as the end time, leave the default values of **now**.

6. To print the report, click **print**.

Save Graphs as GIF Files

Although you can print a report with the Print button, the pagination imposed by browsers may not produce the desired output. As an alternative, you can save each graph in a separate GIF file and then import them into another program for pagination.

> **Note:** This method of saving a graph to a file loses relevant report information such as the class name and time frame. Be sure to note these details associated with each saved graph.

Saving in Microsoft Internet Explorer

To save each graph as a file from a Microsoft Internet Explorer browser window:

1. In the report window, right-click a graph.
2. Click **Save Picture As**.
3. Specify a folder and filename. (The default file type is GIF.)

Saving in Firefox

To save each graph as a file from a Firefox browser window:

1. In the report window, right-click a graph.
2. Click **Save Image As**.
3. Specify a folder and filename. (The default file type is GIF.).

Save Reports as PDF Files

You can make an electronic version of a complete report, by creating an Adobe Acrobat PDF (Portable Document Format) file. You have no control over page breaks, however, so you will probably use this method for one-page reports only.

> **Note:** In order to create PDF files, you need Adobe Acrobat Distiller.

To create a PDF file of a report:

1. In the report window, click **print**.
2. For the printer name, select your Acrobat print driver (such as Acrobat PDFWriter).
3. Click **OK**.
4. Enter a name for the PDF file.
5. Click **Save**.

Display Real-Time Graphs

Use the Statistical Graphing Tool Real-Time graphing feature to view the current real-time bandwidth utilization (in bits per second) for the selected traffic class. A line graph displays the rate at the current time (now) and tracks the rates over the last one thousand seconds (16.7 minutes). The graph updates every second to display real-time data.

To display a real-time graph for a traffic class:

1. Click the **Report** tab.

2. In the Tools section on the left, click **Statistical Graphing**.

> **Note:** In order to use the Statistical Graphing Tool, you must have the Adobe Flash Player plug-in (Version 9). If you do not have it installed, you will be prompted to download and install it before you can proceed.

3. Click the **Real-Time Graphing** tab.

4. In the Class Tree, select the traffic class that you want to graph. As soon as you click on a class name, the corresponding real-time graph begins rendering in the right window pane. **show screen...**

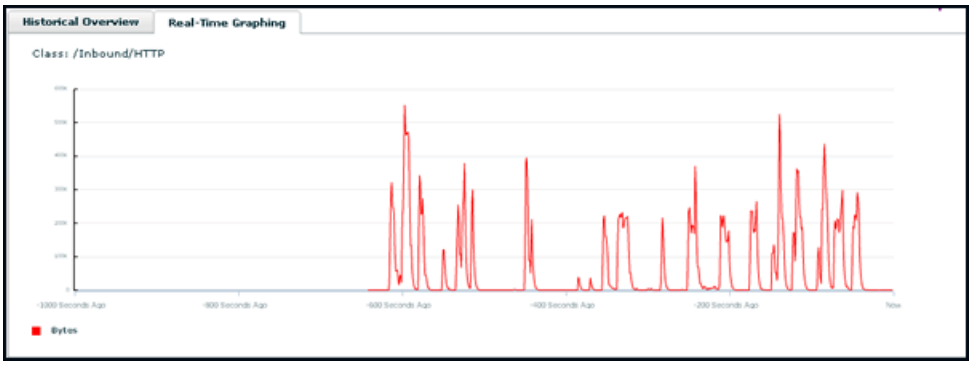

5. Click the ▶ arrow next to the Inbound or Outbound folder to expand the tree. If you have many classes to navigate through, you can filter the class tree by entering a search string in the search field. Only classes containing the string will be displayed.

Additional Resources

For information on analyzing data shown in the Statistical Graphing Tool, see the following PacketGuide Recommendations:

Analyze an Application's Behavior
Identify Performance Saboteurs
Access Wasted Network Capacity

Change the Time Period

By default, the response-time graphs analyze traffic over the last hour, but the time period and end date are user-customizable. The ability to create reports for a specific time period is useful for creating comparison reports. For example, you can print one report that analyzes traffic over the last week (a period of one week, with **now** as the end date and time), and then print a second report that analyzes traffic for the week before that (a period of one week, with a specific date specified for the end date the last day you want to analyze).

To specify a different time period:

1. In the **Show** field, specify the number of minutes, hours, days, weeks, or months for which you want traffic analyzed.

   For example, by entering**2** in the **Show** field and selecting **week** for the unit of time, you create a report with the data PacketShaper collected over the last two weeks. The minimum period you can specify is 1 minute; the maximum is 60 days (or 2 months).

2. To choose a specific end date for the report, select the month and day from the **Before** drop-down lists. If you want the current date as the end date, leave the default values of **now**.

   > **Note:** Although you can select a date from the previous 12 months, the unit may not store data that far back the amount of historical data that is stored depends on how many traffic classes you have in your tree. The unit stores measurement data for a minimum of 31 days in hourly intervals and at least 1 day of data in one-minute intervals.

3. To choose a specific end time for the report, select the time from the drop-down list. If you want the current time as the end time, leave the default values of **now**.

4. Click **update** to update the statistics for the new time period.

   > **Note:** To update the graphs at regular intervals, select an interval from the drop-down list located next to the **update** button.

5. To print the report, click **print**.

Reset RTM Statistics

To reset the response-time statistics:

1. Click the **Monitor** tab.
2. Choose **Response Time** on the **Monitor** drop-down list. The *Monitor Response Time* window appears. **show screen...**

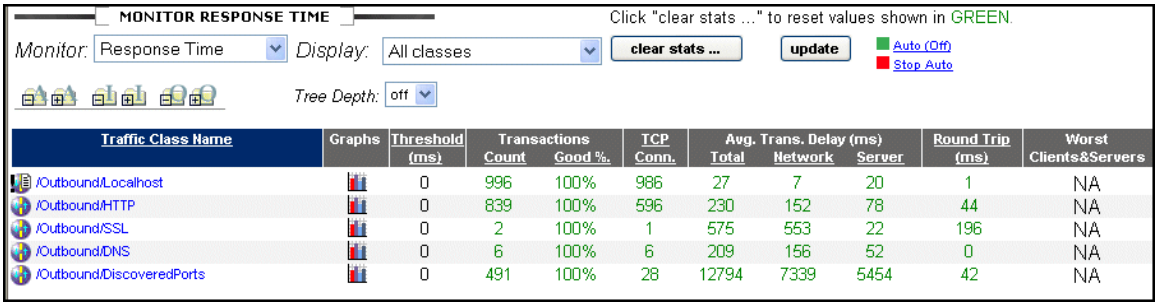

Traffic Class Name	Graphs	Threshold (ms)	Transactions Count	Transactions Good %	TCP Conn.	Avg. Trans. Delay (ms) Total	Avg. Trans. Delay (ms) Network	Avg. Trans. Delay (ms) Server	Round Trip (ms)	Worst Clients&Servers
/Outbound/Localhost		0	996	100%	986	27	7	20	1	NA
/Outbound/HTTP		0	839	100%	596	230	152	78	44	NA
/Outbound/SSL		0	2	100%	1	575	553	22	196	NA
/Outbound/DNS		0	6	100%	6	209	156	52	0	NA
/Outbound/DiscoveredPorts		0	491	100%	28	12794	7339	5454	42	NA

3. Click **clear stats**.

The **clear stats** button zeros-out all response time statistics on the screen, and removes the list of classes as well. Then, as traffic is generated, a new list of classes that have response time statistics available will appear on this screen. Click **update** to see the most current response time data.

Update the Response Time Statistics

The statistics in the *Monitor Response Time* window are captured the moment you choose **Response Time** on the **Monitor** drop-down list. To get a new snapshot, click **update** at the top of the screen. **show screen...**

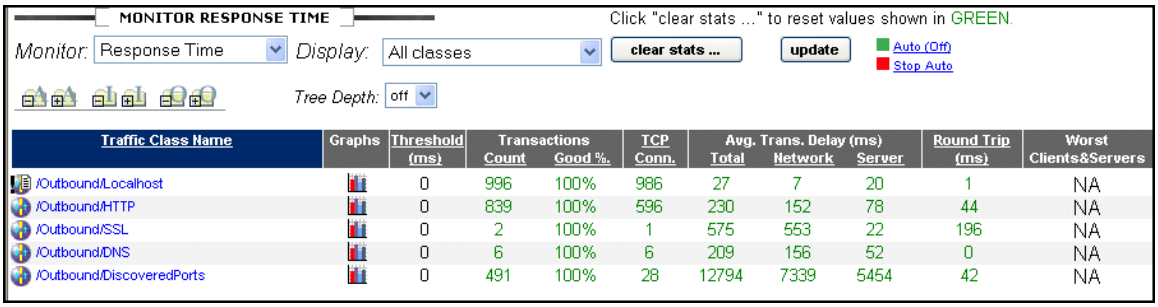

Traffic Class Name	Graphs	Threshold (ms)	Transactions Count	Transactions Good %	TCP Conn.	Avg. Trans. Delay (ms) Total	Avg. Trans. Delay (ms) Network	Avg. Trans. Delay (ms) Server	Round Trip (ms)	Worst Clients&Servers
/Outbound/Localhost		0	996	100%	986	27	7	20	1	NA
/Outbound/HTTP		0	839	100%	596	230	152	78	44	NA
/Outbound/SSL		0	2	100%	1	575	553	22	196	NA
/Outbound/DNS		0	6	100%	6	209	156	52	0	NA
/Outbound/DiscoveredPorts		0	491	100%	28	12794	7339	5454	42	NA

You can set the *Monitor Response Time* window to automatically update at regular intervals. To enable auto-update:

1. Click **Auto**. You are prompted for an interval.
2. Enter the desired number of seconds between updates, and click **OK**.

To stop automatic updating, click **Stop Auto**.

To increase the update interval, do one of the following:

- Click **Auto**. Each time you click **Auto**, three seconds are added to the interval.
- Click **Stop Auto** to disable automatic updates, and then click **Auto** to enter a new interval.

Using this Help System

This help system provides comprehensive information to help you deploy, configure, and use the PacketShaper to control your network.

Find the Information you Need

Table of Contents – Peruse the topics in the help system by collapsing/expanding the hierarchical list in the left pane. You can adjust the width of the Contents pane by dragging its right border. Or click the red arrow (located in the middle of the right border) to hide/display the Contents pane.

Search – Look for a word or a string of words in the topics of the help system. The search results display in the help window; click a topic title to display the topic.

- If you enter multiple words, a list of topics displays that contain *any* of the words in the string.
- To search for topics that contain the exact string, enclose the string in quotation marks.

Related Topics – Links to related topics appear at the end of each topic.

Printing a Topic

1. Click the Expand All button to display any hidden or collapsed text in the topic.
2. Click the Print button. The Print dialog box displays.
3. Select the printer and other options.
4. Click OK to begin printing.

Glossary

M

My Term

My definition