

SGOS Administration Guide

SGOS 6.5.x



Copyright © 2020 Symantec Corp. All rights reserved. Symantec, the Symantec Logo, the Checkmark Logo, Blue Coat, and the Blue Coat logo are trademarks or registered trademarks of Symantec Corp. or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners. This document is provided for informational purposes only and is not intended as advertising. All warranties relating to the information in this document, either express or implied, are disclaimed to the maximum extent allowed by law. The information in this document is subject to change without notice.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE. SYMANTEC CORPORATION PRODUCTS, TECHNICAL SERVICES, AND ANY OTHER TECHNICAL DATA REFERENCED IN THIS DOCUMENT ARE SUBJECT TO U.S. EXPORT CONTROL AND SANCTIONS LAWS, REGULATIONS AND REQUIREMENTS, AND MAY BE SUBJECT TO EXPORT OR IMPORT REGULATIONS IN OTHER COUNTRIES. YOU AGREE TO COMPLY STRICTLY WITH THESE LAWS, REGULATIONS AND REQUIREMENTS, AND ACKNOWLEDGE THAT YOU HAVE THE RESPONSIBILITY TO OBTAIN ANY LICENSES, PERMITS OR OTHER APPROVALS THAT MAY BE REQUIRED IN ORDER TO EXPORT, RE-EXPORT, TRANSFER IN COUNTRY OR IMPORT AFTER DELIVERY TO YOU.

Americas:

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

Rest of the World:

Symantec Limited
Ballycoolin Business Park
Blanchardstown, Dublin 15, Ireland

Email: documentation_inbox@symantec.com

Open source attributions are available in the ProxySG appliance online help. To view the attributions, click **Help** in the appliance to launch the help system, go to the **TOC**, and select **Open Source Attributions for Blue Coat ProxySG**.

Document Number: 231-03113

Document Revision: SGOS 6.5.x—8/2019-D

Contents

Chapter 1: Introduction

Other Documentation	21
Document Conventions	21
Notes and Warnings.....	22
About Procedures	22

Chapter 2: Accessing the Appliance

Accessing the ProxySG Using the Management Console.....	23
About the Management Console Banner.....	25
Viewing the Benefits of Deploying the ProxySG.....	26
Logging Out of the Management Console	31
Accessing the ProxySG Using the CLI.....	32

Section A: Configuring Basic Settings

Configuring the ProxySG Name.....	35
Changing the Login Parameters	35
Changing the Administrator Account Credentials	35
Changing the ProxySG Realm Name	37
Changing the ProxySG Timeout	37
Viewing the Appliance Serial Number	38
Configuring the System Time	38
Synchronizing to the Network Time Protocol.....	40
Appendix: Required Ports, Protocols, and Services	42

Chapter 3: Licensing

License Editions.....	43
Licence Types.....	46
Licensing Terms	47
License Expiration.....	48
Registering and Licensing the Appliance	49
Locating the System Serial Number	49
Obtaining a BlueTouch Online Account.....	50
Registering and Licensing Blue Coat Appliance and Software.....	50
Installing a License on a Registered System.....	51
Manually Installing the License	52
Adding an Add-on License	55
Adding the Add-on License to the ProxySG.....	56
Enabling Automatic License Updates	56

Viewing the Current License Status.....	57
Chapter 4: Controlling Access to the ProxySG	
Moderate Security: Restricting Management Console Access Through the Console Access Control List (ACL)	63
Chapter 5: Backing Up the Configuration	
Section A: About Configuration Archives	
Section B: Archiving Quick Reference	
Archiving Quick Reference Table	72
Section C: Creating and Saving a Standard Configuration Archive	
Section D: Creating and Saving a Secure (Signed) Archive	
Section E: Preparing Archives for Restoration on New Devices	
Creating a Transferable Archive.....	81
Section F: Uploading Archives to a Remote Server	
Creating and Uploading an Archive to a Remote Server	90
Section G: Restoring a Configuration Archive	
Section H: Sharing Configurations	
Section I: Troubleshooting	
Chapter 6: Explicit and Transparent Proxy	
Manually Configure Client Browsers for Explicit Proxy	100
Creating an Explicit Proxy Server with PAC Files	100
Chapter 7: Managing Proxy Services	
Section A: Proxy Services Concepts	
Section B: Configuring a Service to Intercept Traffic	
Changing the State of a Service (Bypass/Intercept)	117
Section C: Creating Custom Proxy Services	
Section D: Proxy Service Maintenance Tasks	
Section E: Global Options for Proxy Services	
Proxy Service Global Options	130
Managing Licensed User Connection Limits (ProxySG to Server)	136
About User Limits	136
Tasks for Managing User Limits	138
Viewing Concurrent Users	141
Section F: Exempting Requests From Specific Clients	
Adding Static Bypass Entries	143
Section G: Trial or Troubleshooting: Restricting Interception From Clients or To Servers	

Restricted Intercept Topics	147
Section H: Reference: Proxy Services, Proxy Configurations, and Policy	
Chapter 8: Intercepting and Optimizing HTTP Traffic	
Section A: About the HTTP Proxy	
Section B: Changing the External HTTP (Transparent) Proxy Service to Intercept All IP Addresses on Port 80	
Section C: Managing the HTTP Proxy Performance	
About HTTP Compression	162
Understand Compression Behavior	163
Compression Exceptions	164
Configuring Compression	164
Notes	167
About the HTTP Object Caching Policy Global Defaults	167
Setting the HTTP Default Object Caching Policy	171
Section D: Selecting an HTTP Proxy Acceleration Profile	
Configuring the HTTP Proxy Profile	180
Section E: Using a Caching Service	
Prerequisite for Using CachePulse	182
Enabling CachePulse	182
Downloading the CachePulse Database	183
Section F: Fine-Tuning Bandwidth Gain	
Allocating Bandwidth to Refresh Objects in Cache	184
Section G: Caching Authenticated Data (CAD) and Caching Proxy Authenticated Data (CPAD)	
Section H: Viewing HTTP/FTP Statistics	
Viewing the Number of HTTP/HTTPS/FTP Objects Served	202
Viewing the Number of HTTP/HTTPS/FTP Bytes Served	203
Viewing Active Client Connections	204
Viewing HTTP/HTTPS/FTP Client and Server Compression Gain Statistics	204
Viewing HTTP/FTP Client Compressed Gain Statistics	205
Viewing HTTP/FTP Server Compressed Gain Statistics	205
Section I: Supporting IWA Authentication in an Explicit HTTP Proxy	
Section J: Supporting Authentication on an Upstream Explicit Proxy	
Deployment Scenarios	209
Section K: Detect and Handle WebSocket Traffic	
How the ProxySG Appliance Handles an Upgrade Request	210
Feature Limitations	211

Chapter 9: Managing the SSL Proxy

IPv6 Support	214
Working with SSL Traffic.....	215

Section A: Intercepting HTTPS Traffic

Configuring the SSL Proxy in Explicit Proxy Mode	220
Specifying an Issuer Keyring and CCL Lists for SSL Interception	220
Using Client Consent Certificates.....	221
Downloading an Issuer Certificate	221
Warn Users When Accessing Websites with Untrusted Certificates	225
Presenting Untrusted Certificates to a Browser.....	225
Set the Behavior when Encountering Untrusted Certificates	225

Section B: Configuring SSL Rules through Policy

Section C: Viewing SSL Statistics

Viewing SSL History Statistics.....	232
Unintercepted SSL Data	232
Unintercepted SSL Clients	233
Unintercepted SSL Bytes	233

Section D: Using STunnel

Configuring STunnel.....	235
Viewing STunnel Results.....	237
Application Mix.....	239
Viewing Session Statistics	239
Viewing Protocol Details.....	240
Access Logging	240

Section E: Tapping Decrypted Data with Encrypted Tap

Viewing Encrypted Tap Results.....	243

Section F: Working with an HSM Appliance

Working with the SafeNet Java HSM	244
Before You Begin	244
Add an HSM	245
Add an HSM Keyring.....	245
Adding an HSM Keygroup.....	246
Write HSM Policy	247

Section G: Advanced Topics

Chapter 10: Accelerating File Sharing

Configuring the ProxySG CIFS Proxy	258
About Windows Security Signatures	258
Intercepting CIFS Services	261

Configuring SMBv1 Options	262
Configuring SMBv2 Options	267
Enabling CIFS Access Logging.....	268
Reviewing CIFS Protocol Statistics.....	268
Chapter 11: Managing Outlook365 Applications	
Section A: The Outlook Proxies	
Section B: Endpoint Mapper and MAPI Configuration	
Reviewing Endpoint Mapper Proxy Statistics	279
Configuring the MAPI Proxy	280
Reviewing MAPI Statistics.....	281
Join the Branch Peer to the Primary Domain	285
Section C: Intercept Skype for Business	
Configure ProxySG for for Skype and Lync Interception	289
Chapter 12: Managing the File Transport Protocol (FTP) Proxy	
Configuring the ProxySG for Native FTP Proxy	295
Modifying the FTP Proxy Service	295
Configuring the FTP Proxy.....	296
Configuring FTP Clients for Explicit Proxy	297
Chapter 13: Managing the Domain Name Service (DNS) Proxy	
Chapter 14: Managing a SOCKS Proxy	
IPv6 Support	306
Configuring the SOCKS Proxy	307
Viewing SOCKS History Statistics	308
Viewing SOCKS Clients	308
Viewing SOCKS Connections.....	309
Viewing SOCKS Client and Server Compression Gain Statistics	309
Chapter 15: Managing Shell Proxies	
Configuring the Telnet Shell Proxy Service Options.....	315
Viewing Shell History Statistics.....	317
Chapter 16: Configuring and Managing an HTTPS Reverse Proxy	
Section A: About the HTTPS Reverse Proxy	
Section B: Configuring the HTTPS Reverse Proxy	
About Mutual SSL Authentication	325
Section C: Configuring HTTP or HTTPS Origination to the Origin Content Server	

Chapter 17: Using the ProxySG in an IPv6 Environment

Using the ProxySG in an ISATAP Network	334
IPv6 Support on the ProxySG	337
Configuring an ADN for an IPv6 Environment.....	346
Optimizing ISATAP Traffic.....	347
Configuring IPv6 Global Settings.....	349

Chapter 18: Client Geolocation

Prerequisites for Using Geolocation	352
Enable Geolocation	353
Download the Geolocation Database	354
Determine Locations of IP Addresses for Incoming Connections.....	356
Troubleshoot Geolocation	358
Access Log Errors	359
Remove Geolocation Settings	360

Chapter 19: Web Application Protection

Section A: Using Application Protection

Enabling Application Protection	362
Testing the Application Protections.....	363
Verifying the Database Download.....	364

Section B: Understanding the Risk Score

Creating Actions Based on the Client's Risk Score	365

Section C: Creating an Application Protection Solution

About Injection Attacks	367
Preventing an SQL Injection Attack.....	368
Configuring Injection Protection	369
Preventing a Null-Byte Injection Attack.....	369
Preventing an Invalid Multipart Form Attack.....	370
Preventing an HTTP Parameter Pollution Attack.....	370
Preventing a Multiple Encoding Attack	371
Creating Custom Detection Rules	372

Section D: Advanced Features for Web Application Protection

Section E: Reference Information

Chapter 20: Filtering Web Content

Section A: Web Content Filtering Concepts

About Application Filtering	379
Web Content Filtering Process Flow	380
About Symantec WebFilter and the WebPulse Service.....	381
About Dynamic Categorization	382

Considerations Before Configuring WebPulse Services.....	387
Section B: Setting up a Web Content Filter	
Enabling a Content Filter Provider	392
Downloading the Content Filter Database	394
About Database Updates	394
Downloading a Content Filter Database	395
Viewing the Status of a Database Download.....	396
Expiry Date for the Database.....	397
Viewing the Available Categories or Testing the Category for a URL	397
Testing the Application and Operation for a URL	397
Section C: Configuring Symantec WebFilter and WebPulse	
Disabling Dynamic Categorization	399
Specifying a Custom Time Period to Update Symantec WebFilter	400
Section D: Configuring a Local Database	
Local Database Matching Example	406
Selecting and Downloading the Local Database.....	407
Section E: Configuring Internet Watch Foundation	
Section F: Configuring a Third-Party Vendor	
Section G: About Blue Coat Categories for YouTube	
Setting the YouTube Server Key	413
Distinguishing Blue Coat Categories for YouTube in the Access Log	413
Section H: Applying Policy	
Policy Examples Using the Application Control Objects.....	432
Section I: Troubleshooting	
Chapter 21: Configuring Threat Protection	
Adding an ICAP service for Content Scanning.....	447
Chapter 22: Malicious Content Scanning Services	
Section A: About Content Scanning	
Section B: Configuring ICAP Services	
Creating an ICAP Service	474
Managing ICAP Health Checks	479
Monitoring ICAP Health Metrics	480
Configuring ICAP Feedback	482
Customizing ICAP Patience Text	483
HTTP Patience Text	484
FTP Patience Text	487

Section C: Securing Access to an ICAP Server

Using Secure ICAP	488
Using a Crossover Cable.....	491
Configuring ICAP Using a Crossover Cable.....	492
Using a Private Network	493

Section D: Monitoring ICAP Requests and Sessions

Introduction to ICAP Request Monitoring	496

Section E: Creating ICAP Policy

Using ICAP Headers in Policy	514

Section F: Managing Virus Scanning

Chapter 23: Configuring Service Groups

Creating a Service Group.....	523

Chapter 24: Managing Streaming Media

Section A: Concepts: Streaming Media

Limitation	533
About Microsoft Smooth Streaming.....	533
About Adobe HDS	534
About Windows Media	534
About Processing Streaming Media Content	538
Limiting Bandwidth	540
Caching Behavior: Proxy Specific	542
Caching Behavior: Video-on-Demand	543
Splitting Behavior: Live Broadcast	543
About Streaming Media Authentication	547
Apple HLS Authentication	549

Section B: Configuring Streaming Media

Configuring the HTTP Streaming Proxy	550
Configuring the Windows Media, Real Media, and QuickTime Proxies.....	554
Limiting Bandwidth	556
Configuring Bandwidth Limitation—Fast Start (WM).....	557
Limiting Bandwidth for Smooth Streaming.....	557
Legacy Streaming Log Format	559
Reporter Streaming Log Format	561
Viewing Streaming History Statistics	562
Viewing Current and Total Streaming Data Statistics.....	562

Section C: Additional Windows Media Configuration Tasks

Section D: Configuring Windows Media Player

Section E: Configuring RealPlayer

Section F: Configuring QuickTime Player

Section G: Using the Flash Streaming Proxy

Configuring the Flash Streaming Proxy	584
Configuring Client Browsers for Explicit Proxy	584
Intercepting the RTMP Service (Transparent Deployment)	585
Intercepting the Explicit HTTP Service (Explicit Deployment)	585
When VOD Content Gets Cached.....	588
Proxy Chaining.....	588
CDN Interoperability Support	589

Section H: Supported Streaming Media Clients and Protocols

Chapter 25: Bandwidth Management

Configuring Bandwidth Allocation	602
Bandwidth Management Statistics.....	604
Current Class Statistics.....	604
Total Class Statistics.....	605
Using Policy to Manage Bandwidth	606

Chapter 26: Configuring Access Logging

Configuring a Log for Uploading	620
Viewing Access-Log Statistics.....	623
Viewing the Access Log Tail.....	623
Viewing the Log File Size.....	624
Viewing Access Logging Status	625

Chapter 27: Configuring the Upload Client

Importing an External Certificate.....	630
Deleting an External Certificate	631
Digitally Signing Access Logs.....	631
Introduction to Digitally Signing Access Logs	632
Configuring the Upload Client to Digitally Sign Access Logs	632

Chapter 28: Creating and Editing an Access Log Facility

Creating a Log Facility	641
Editing an Existing Log Facility.....	643
Associating a Log Facility with a Protocol.....	644
Configuring Global Settings.....	646

Chapter 29: Creating Custom Access Log Formats

Creating a Custom or ELFF Log Format	652
Creating Custom Log Formats Based on Reserved Log Formats	653

Chapter 30: Access Log Formats

Action Field Values	661

Chapter 31: Statistics

Viewing Bandwidth Details for Proxies or Services	668
Viewing Traffic Distribution	670
Viewing Per-Proxy or Per-Service Statistics.....	671
About Bypassed Bytes	671
About the Default Service Statistics	672
Viewing NetFlow Statistics	672
Viewing Traffic History	673
Supported Proxies and Services	675
Viewing the Application Mix Report.....	676
Viewing Bandwidth Details for Web Applications.....	678
Viewing the Application History Report	681
Viewing System Statistics	682
Resources Statistics	682
Contents Statistics	686
Event Logging Statistics	689
Failover Statistics.....	690
Active Sessions—Viewing Per-Connection Statistics.....	690
Example Scenarios Using Active Sessions for Troubleshooting	691
Analyzing Proxied Sessions.....	691
Analyzing Bypassed Connections Statistics.....	703
Viewing Errored Sessions and Connections	705

Chapter 32: Configuring an Application Delivery Network**Section A: ADN Overview**

ADN Modes.....	715
Multiple Concentrators in a Transparent ADN Deployment.....	716
Discovery of Upstream Concentrators.....	718

Section B: Configuring an ADN

Introduction to Configuring an ADN.....	724
Enabling Explicit ADN Connections	729
Advertising Server Subnets	729
Configuring the Tunnel Mode	730
Configuring IP Address Reflection	735
Enabling ProxyClient Support	737

Section C: Securing the ADN

Securing a Managed ADN	740
Enabling Device Authentication	740

Configuring Connection Security	742
Enabling Device Authorization.....	743
Section D: Configuring Load Balancing	
Introduction to Load Balancing	746
Section E: Configuring Advanced ADN Settings	
Configuring an ADN Node as an Internet Gateway	750
Configuring the Byte-Cache Dictionary Size	752
Manually Resizing the Byte Cache Dictionaries From the Statistics Tab.....	753
Manually Resizing Byte Cache Dictionaries from the Byte Caching Tab	754
Section F: Monitoring the ADN	
Reviewing ADN History	758
Reviewing ADN Active Sessions	759
Monitoring Adaptive Compression	760
Section G: Related CLI Syntax to Configure an ADN	
Section H: Policy	
Section I: Troubleshooting	
Chapter 33: WCCP Configuration	
Configuring WCCP on the ProxySG.....	780
Creating the WCCP Configuration on the ProxySG	780
Modifying the WCCP Configuration	785
Disabling WCCP.....	786
Viewing WCCP Statistics and Service Group Status.....	786
Chapter 34: TCP/IP Configuration	
PMTU Discovery.....	791
Chapter 35: Routing on the ProxySG	
Distributing Traffic Through Multiple Default Gateways	796
ProxySG Specifics.....	796
Switching to a Secondary Default Gateway	797
Routing in Transparent Deployments	798
Outbound Routing.....	798
Inbound Routing	799
About Trust Destination MAC.....	799
About Static Routes.....	800
Using Return-to-Sender (RTS).....	801
Chapter 36: Configuring Failover	
Configuring Failover Groups.....	808
Viewing Failover Statistics	810

Chapter 37: Configuring DNS

Adding DNS Servers to the Primary or Alternate Group	816
Resolving Hostnames Using Name Imputing Suffixes.....	819
Adding and Editing DNS Name Imputing Suffixes	820
Changing the Order of DNS Name Imputing Suffixes	821

Chapter 38: Virtual IP Addresses

Creating a VIP	823
Deleting a VIP	824

Chapter 39: Configuring Private Networks

Configuring Private Subnets.....	826
Configuring Private Domains.....	827

Chapter 40: Managing Routing Information Protocols (RIP)

Installing RIP Configuration Files.....	831

Chapter 41: SOCKS Gateway Configuration

Section A: Configuring a SOCKS Gateway

Adding a SOCKS Gateway.....	840
Creating SOCKS Gateway Groups.....	843
Configuring Global SOCKS Defaults.....	845
Configuring the SOCKS Gateway Default Sequence	847

Section B: Using SOCKS Gateways Directives with Installable Lists

Creating a SOCKS Gateway Installable List	854

Chapter 42: TCP Connection Forwarding

Configuring TCP Connection Forwarding	862
Copying Peers to Another ProxySG in the Cluster	863
Removing a Peer.....	864

Chapter 43: Configuring the Upstream Network Environment

Section A: Overview

Section B: About Forwarding

Section C: Configuring Forwarding

Creating Forwarding Hosts and Groups.....	875
Creating Forwarding Hosts	875
Creating Forwarding Groups	877
Configuring Global Forwarding Defaults.....	879
Configuring the Forwarding Default Sequence	881

Section D: Using Forwarding Directives to Create an Installable List

Creating a Forwarding Installable List.....	889

Chapter 44: Using Policy to Manage Forwarding

Chapter 45: About Security

Controlling User Access with Identity-based Access Controls	898

Chapter 46: Controlling Access to the Internet and Intranet

Section A: Managing Users

Viewing Logged-In Users.....	900

Section B: Using Authentication and Proxies

About Authentication Modes	908
Setting the Default Authenticate Mode Property.....	910
About Origin-Style Redirection	910
Selecting an Appropriate Surrogate Credential	911
Manually Entering Top-Level Domains (TLDs).....	911
Configuring Transparent Proxy Authentication	911
Permitting Users to Log in with Authentication or Authorization Failures	912
Using Guest Authentication	914
Using Default Groups.....	915

Section C: Using SSL with Authentication and Authorization Services

Section D: Creating a Proxy Layer to Manage Proxy Operations

Section E: Forwarding BASIC Credentials

Chapter 47: Local Realm Authentication and Authorization

Creating a Local Realm	933
Changing Local Realm Properties.....	934

Chapter 48: CA eTrust SiteMinder Authentication

Creating a SiteMinder Realm	949
Configuring SiteMinder Agents.....	949
Configuring SiteMinder Servers.....	951
Defining SiteMinder Server General Properties.....	952
Configuring Authorization Settings for SiteMinder	954
Configuring General Settings for SiteMinder	955

Chapter 49: Certificate Realm Authentication

Configuring Certificate Realms	960
Creating a Certificate Realm.....	960
Configuring Certificate Realm Properties	961
Defining General Certificate Realm Properties	963
Specifying an Authorization Realm.....	964

Chapter 50: Oracle COREid Authentication

Creating a COREid Realm	974
Configuring Agents for COREid Authentication	975
Configuring the COREid Access Server	976
Configuring the General COREid Settings	977

Chapter 51: SAML Authentication

About SAML.....	982
Federation and Metadata	982
Assertions	982
Profiles and Bindings.....	983
Requirements for SAML Authentication	983
Checklist: Preparing the IDP	983
An Overview of the Authentication Process	984
Set up SAML Authentication	986
Export the IDP Metadata File.....	986
Prepare the Appliance.....	988
Configure the CCL	988
Create an HTTPS Reverse Proxy Service	988
Configure SAML Attributes	989
Configure General Settings for SAML	989
Create the SAML Realm	990
Configure SAML Authorization.....	992
Policy Conditions	993
Configure the IDP.....	993
Configure AD FS	993
Configure SiteMinder	996
Configure Oracle	1000
Configure Shibboleth.....	1003
Prevent Dropped Connections When Policy is Set to Deny.....	1004
Backing Up ProxySG Configuration: Considerations for SAML.....	1004
Save Keyrings Before Backing up the Configuration	1005
Import Saved Keyrings After Restoring the Configuration.....	1005
Re-Import the ProxySG Certificate to the Trust List (AD FS).....	1005

Chapter 52: Integrating the Appliance with Your Windows Domain

Integrate the ProxySG Appliance into the Windows Domain.....	1007
Join the ProxySG Appliance to the Windows Domain.....	1008
Edit a Windows Domain	1009
Configure SNMP Traps for the Windows Domain	1011

Chapter 53: Integrating ProxySG Authentication with Active Directory Using IWA

About IWA Challenge Protocols.....	1014
About IWA Failover.....	1014
Preparing for a Kerberos Deployment	1016
Enabling Kerberos in an IWA Direct Deployment.....	1016
Enabling Kerberos in a BCAAA Deployment.....	1017
Configuring IWA on the ProxySG Appliance	1018
Creating an IWA Realm	1018
Configuring IWA Servers	1020
Defining IWA Realm General Properties	1025
Creating the IWA Authentication and Authorization Policies.....	1027
Creating an IWA Authentication Policy	1028
Creating a Guest Authentication Policy	1030
Creating an IWA Authorization Policy.....	1031
Configuring Client Systems for Single Sign-On.....	1033
Configure Internet Explorer for Single Sign-On.....	1033
Configure Firefox for Single Sign-On.....	1034
Using IWA Direct in an Explicit Kerberos Load Balancing/Failover Scenario.....	1034

Chapter 54: Kerberos Constrained Delegation

Chapter 55: LDAP Realm Authentication and Authorization

Creating an LDAP Realm on the ProxySG.....	1046
About LDAP Realms	1046
Creating an LDAP Realm.....	1047
Configuring LDAP Properties on the ProxySG.....	1048
Configuring LDAP Servers.....	1048
Defining LDAP Base Distinguished Names.....	1050
Defining LDAP Search & Group Properties	1052
Customizing LDAP Objectclass Attribute Values.....	1056
Defining LDAP General Realm Properties.....	1057

Chapter 56: Novell Single Sign-on Authentication and Authorization

Creating a Novell SSO Realm	1067
Novell SSO Agents	1068
Adding LDAP Servers to Search and Monitor for Novell SSO	1069
Querying the LDAP Novell SSO Search Realm	1070
Configuring Authorization	1071
Defining Novell SSO Realm General Properties.....	1072

Chapter 57: Policy Substitution Realm

Creating a Policy Substitution Realm	1080

Configuring User Information	1081
Creating a List of Users to Ignore.....	1083
Configuring Authorization	1083
Defining Policy Substitution Realm General Properties	1084
Creating the Policy Substitution Policy	1086
Chapter 58: RADIUS Realm Authentication and Authorization	
Creating a RADIUS Realm	1090
Defining RADIUS Realm Properties.....	1090
Defining RADIUS Realm General Properties.....	1092
Chapter 59: Configuring the ProxySG as a Session Monitor	
Chapter 60: Sequence Realm Authentication	
Creating a Sequence Realm	1108
Adding Realms to a Sequence Realm	1109
Defining Sequence Realm General Properties	1111
Chapter 61: Managing X.509 Certificates	
Section A: PKI Concepts	
Certificate Chains	1115
Server-Gated Cryptography and International Step-Up	1117
Section B: Using Keyrings and SSL Certificates	
Creating a Keyring.....	1119
Deleting an Existing Keyring and Certificate	1122
Providing Client Certificates in Policy	1123
Add Certificates to the ProxySG Appliance	1123
Create a Keydata File.....	1124
Import Certificates onto the ProxySG Appliance.....	1125
Group Related Client Keyrings into a Keylist	1125
Specify the Client Certificates to be Used in Policy	1127
Specify the Client Certificates to be Used in Policy in the VPM	1127
Specify the Client Certificates to be Used in Policy in CPL	1128
Section C: Managing Certificates	
Managing SSL Certificates.....	1132
Creating Self-Signed SSL Certificates.....	1133
Importing a Server Certificate	1134
Using Certificate Revocation Lists	1134
Section D: Using External Certificates	
Section E: Advanced Configuration	
Managing CA Certificate Lists.....	1143

Creating a CA Certificate List:	1144
Updating a CA Certificate List.....	1146
Configuring Download of CCL Updates from Symantec	1146
Managing Cached Intermediate Certificates	1148
Turn off Intermediate Certificate Caching	1149
View Cached Intermediate Certificates	1149
Clear Cached Intermediate Certificates	1150
Section F: Checking Certificate Revocation Status in Real Time (OCSP)	
Creating and Configuring an OCSP Responder	1155
Chapter 62: Managing SSL Traffic	
Section A: SSL Client Profiles	
Editing an SSL Client	1166
Associating a Keyring, Protocol, and CCL with the SSL Client	1166
Changing the Cipher Suite of the SSL Client	1167
Section B: SSL Device Profiles	
Section C: Notes and Troubleshooting	
Chapter 63: Windows Single Sign-on Authentication	
Creating a Windows SSO Realm	1178
Configuring Windows SSO Agents	1178
Configuring Windows SSO Authorization	1180
Defining Windows SSO Realm General Properties	1182
Chapter 64: Using XML Realms	
Creating an XML Realm	1190
Configuring XML Servers.....	1191
Configuring XML Options	1192
Configuring XML Realm Authorization	1193
Configuring XML General Realm Properties	1194
Chapter 65: Forms-Based Authentication	
Creating and Editing a Form.....	1204
Setting Storage Options	1206
Chapter 66: Authentication and Authorization Errors	
Chapter 67: Configuring Adapters and Virtual LANs	
Changing the Default Adapter and Interface Settings	1237
About Multiple IP Addresses	1237
Configuring a Network Adapter	1237
Improve Resiliency or Create a Bigger Pipe with an Aggregate Interface	1243

Viewing Interface Statistics	1248
Chapter 68: Software and Hardware Bridges	
Configuring a Software Bridge	1255
Chapter 69: Configuring Management Services	
Creating a Management Service	1268
Creating a Notice and Consent Banner for the Management Console	1273
Managing the SSH Console	1273
Managing the SSH Host Key Pairs	1274
Creating a Notice and Consent Banner for SSH	1275
Managing SSH Client Keys	1275
Chapter 70: Preventing Denial of Service Attacks	
Creating the CPL	1287
Chapter 71: Authenticating a ProxySG	
Obtaining a ProxySG Appliance Certificate	1292
Automatically Obtaining an Appliance Certificate	1292
Manually Obtaining an Appliance Certificate	1292
Creating an SSL Device Profile for Device Authentication	1296
Chapter 72: Monitoring the ProxySG	
Section A: Using Director to Manage ProxySG Systems	
Automatically Registering the ProxySG with Director	1300
Registration Requirements	1301
Registering the ProxySG with Director	1301
Section B: Monitoring the System and Disks	
System Configuration Summary	1304
Viewing System Environment Sensors	1305
Viewing Disk Status and Taking Disks Offline	1306
Viewing SSL Accelerator Card Information	1307
Section C: Configuring Event Logging and Notification	
Selecting Which Events to View	1309
Setting Event Log Size	1310
Enabling Event Notification	1310
Syslog Event Monitoring	1312
Securely Retrieving Event Logs from the Appliance	1314
Viewing Event Log Configuration and Content	1316
Section D: Monitoring Network Devices (SNMP)	
Configuring SNMP Communities	1325
Configuring SNMP for SNMPv1 and SNMPv2c	1326

Adding Community Strings for SNMPv1 and SNMPv2c.....	1327
Configuring SNMP Traps for SNMPv1 and SNMPv2c.....	1328
Configuring SNMP for SNMPv3.....	1330
About Passphrases and Localized Keys	1330
Configuring SNMP Users for SNMPv3	1330
Configuring SNMP Traps and Informs for SNMPv3	1333
Section E: Configuring Health Monitoring	
About the Health Monitoring Metric Types.....	1338
About the General Metrics.....	1340
About the Licensing Metrics.....	1341
About the Status Metrics.....	1343
Snapshot of the Default Threshold Values and States.....	1345
Changing Threshold and Notification Properties	1347
Viewing Health Monitoring Statistics.....	1349
Chapter 73: Verifying the Health of Services Configured on the ProxySG	
Section A: Overview	
Background DNS Resolution.....	1356
Section B: About Blue Coat Health Check Components	
Health Check Tests.....	1359
Section C: Configuring Global Defaults	
Changing Health Check Default Settings	1366
Configuring Health Check Notifications	1369
Section D: Forwarding Host and SOCKS Gateways Health Checks	
Section E: DNS Server Health Checks	
Section F: Authentication Health Checks	
Section G: Virus Scanning and Content Filtering Health Checks	
Section H: Managing User-Defined Health Checks	
Section I: Viewing Health Check Statistics	
Health Check Topics	1392
About Health Check Statistics	1392
Section J: Using Policy	
Chapter 74: Maintaining the ProxySG	
Performing Maintenance Tasks.....	1399
Upgrading the ProxySG Appliance	1404
Managing ProxySG Systems.....	1404
Setting the Default Boot System.....	1406
Locking and Unlocking ProxySG Systems.....	1406

Replacing a ProxySG System.....	1407
Deleting a ProxySG System	1407

Chapter 75: Diagnostics

Diagnostic Reporting (Service Information).....	1412
Sending Service Information Automatically	1412
Managing the Bandwidth for Service Information	1413
Configure Service Information Settings.....	1414
Creating and Editing Snapshot Jobs.....	1417
Packet Capturing (PCAP—the Job Utility)	1419
PCAP File Size	1419
PCAP File Name Format	1420
Common PCAP Filter Expressions.....	1420
Configuring Packet Capturing	1421
Core Image Restart Options	1425
Diagnostics: Symantec Customer Experience Program and Monitoring	1426

Chapter 76: XML Protocol

- Section A: Authenticate Request**
- Section B: Authenticate Response**
- Section C: Authorize Request**
- Section D: Authorize Response**

Chapter 1: Introduction

This audience for this document is network administrators who are responsible for managing Blue Coat® ProxySG® appliances. This document provides reference information and procedures to configure SGOS™ version 6.5, and includes topics for Application Delivery Network (ADN), including Application Acceleration and Secure Web Gateway solutions.

The information in this document supersedes information in the ProxySG Management Console Online Help System.

Other Documentation

Other documentation for the 6.5.x software line is available:

- ❑ SGOS 6.5.x Release Notes
- ❑ SGOS Upgrade/Downgrade Guide
- ❑ Command Line Interface Reference
- ❑ Visual Policy Manager Reference (includes some advanced policy tasks)
- ❑ Content Policy Language Reference

Document Conventions

The following table lists the typographical and Command Line Interface (CLI) syntax conventions used in this manual.

Table 1–1 Document Conventions

Conventions	Definition
<i>Italics</i>	The first use of a new or Blue Coat-proprietary term.
<code>Courier font</code>	Screen output. For example, command line text, file names, and Content Policy Language (CPL).
<code><i>Courier Italics</i></code>	A command line variable that is to be substituted with a literal name or value pertaining to the appropriate facet of your network system.
Courier Boldface	A Blue Coat literal to be entered as shown.
Arial Boldface	Screen elements in the Management Console.
{ }	One of the parameters enclosed within the braces must be supplied
[]	An optional parameter or parameters.

Table 1–1 Document Conventions (Continued)

	Either the parameter before or after the pipe character can or must be selected, but not both.

Notes and Warnings

The following is provided for your information and to caution you against actions that can result in data loss or personal injury:

Note: Supplemental information that requires extra attention.

Important: Critical information that is not related to equipment damage or personal injury (for example, data loss).

WARNING! Used *only* to inform you of danger of personal injury or physical damage to equipment. An example is a warning against electrostatic discharge (ESD) when installing equipment.

About Procedures

Many of the procedures in this guide begin:

- ❑ **Select Configuration > *TabName***, if you are working in the Management Console, or
- ❑ **From the (config) prompt**, if you are working in the command line interface (CLI).

Blue Coat assumes that you are logged into the first page of the Management Console or entered into configuration mode in the CLI.

In most cases, procedures in this guide tell you how to perform a task in the Management Console, even if there is a CLI equivalent.

Chapter 2: Accessing the Appliance

This section provides procedures for accessing the ProxySG so that you can perform administrative tasks using the Management Console and/or the command-line interface. It assumes that you have performed the first-time setup using the Serial Console or the front panel and that you have minimally specified an IP address, IP subnet mask, IP gateway, and DNS server, and that you have tested the appliance and know that it is up and running on the network. If you have not yet done this, refer to the [hardware guides](#) for your appliance model.

This section includes the following topics:

- ❑ ["Accessing the ProxySG Using the Management Console"](#) on page 23
- ❑ ["Accessing the ProxySG Using the CLI"](#) on page 32
- ❑ ["Configuring Basic Settings"](#) on page 34


Accessing the ProxySG Using the Management Console

The Management Console is a graphical Web interface that allows you to manage, configure, monitor, and upgrade the ProxySG from any location. To determine the browser and Java requirements for the Management Console, refer to the *SGOS Release Notes*.

Note: When you access the Management Console home page, if you see a `host mismatch` or an `invalid certificate` message, you must recreate the security certificate used by the HTTPS-Console. For information on changing the security certificate, see ["Managing the HTTPS Console \(Secure Console\)"](#) on page 1270.


To log in to the Management Console:


1. In the Web browser's address bar, enter `https://
appliance_IP_address:8082`
The default management port is 8082.
For example, if the IP address configured during first-time installation is 192.168.0.6, type `https://192.168.0.6:8082` in the Web browser.



2. Enter the user name and password that you created during first-time set up. The Management Console **Statistics > Summary > Efficiency** page displays. For information on the details displayed on the **Statistics > Summary** tab, see ["Viewing Efficiency and Performance Metrics"](#) on page 26 and ["Monitoring System Resources and Connectivity Metrics"](#) on page 28.


Note: All successful and failed logon attempts are recorded in the event log.





About the Management Console Banner

The Management Console banner displays across the top of the Web browser, after you have logged in to the ProxySG.



The Management Console banner provides the following information:

- ❑ Appliance identification—the appliance name, hardware model, hardware serial number, and the software version.
- ❑ Appliance health status—The health state is represented by a text string and a color that corresponds to the health of the system (OK-green, Warning- yellow or Critical -red). The system health changes when one or more of the health metrics reaches a specified threshold or returns to normal. The health state indicator is polled and updated every 10 seconds on the ProxySG.

To obtain more information about the health state, click the **Health:** status link — Ok, Warning, Critical. The **Statistics > Health** page displays; it lists the current condition of the system's health monitoring metrics. See ["Verifying the Health of Services Configured on the ProxySG"](#) on page 1353 for more information about the health monitoring metrics.

- ❑ License status and version—Your ProxySG license includes all the component licenses for the Proxy features that you have purchased. To view a list of the license components and their expiration date, go to the **Maintenance > Licensing > View** tab.

By default, for a new ProxySG appliance, the trial edition is enabled— at initial set-up you had elected to use either the Proxy edition or the MACH5 edition. For the first 60 days of the trial period, all licensable components for the edition you chose are active and available to use. During the trial period, the Base SGOS license allows unlimited concurrent users. To view the specifics of your trial edition license, click the **Trial Period** link.

- ❑ Symantec product documentation and customer support links. You must have a Blue Touch Online account to access documentation and to request support. To log out of the ProxySG Management Console, click the **Log Out** link.

Viewing the Benefits of Deploying the ProxySG

The **Statistics > Summary** page displays the role of the ProxySG in boosting the performance of traffic within your network using its acceleration, optimization, policy control, and caching techniques. The **Summary** page visually demonstrates the overall performance and efficiency of your network.

If you have just completed initial setup and have not configured the ProxySG to intercept any traffic, the **Summary** page will not display much information. For example, you cannot view bandwidth efficiency and savings for traffic being intercepted by the ProxySG.


Note: To view performance statistics, retrieve your license and create/enable services on the ProxySG. For information on enabling services, see [Chapter 7: "Managing Proxy Services"](#) on page 109. For licensing details, see [Chapter 3: "Licensing"](#) on page 43.

When the ProxySG is deployed and configured to meet your business needs, the **Summary** page monitors and reports information on your network traffic and applications. The on-screen information is automatically refreshed every 60 seconds.

Viewing Efficiency and Performance Metrics

The **Statistics > Summary > Efficiency** tab displays the bandwidth gain achieved within your network in the Savings panel, and the performance of each interface in the Interface Utilization panel on the ProxySG. These metrics represent the last hour of traffic on the ProxySG, and are updated every 60 seconds.

The Savings panel displays the top 5 services that are intercepted by the ProxySG, in your network. For detailed information on each service, click the service and view the details in the **Statistics > Traffic History** page.



- ❑ **Service:** A service represents the type of traffic that is being intercepted; the top 5 services are ranked in descending order of bytes saved.
- ❑ **Bytes Saved Last Hour:** Bytes saved display bandwidth savings in the last 60 minutes. It represents data that did not traverse the WAN because of object and byte caching, protocol optimization, and compression. It is calculated as:

$$\text{Client Bytes} - \text{Server Bytes}$$
 where *Client Bytes* is the data rate calculated to and from the client on the client-side connection, and *Server Bytes* is the data rate calculated to and from the server on the server-side connection.

For Inbound ADN, bytes saved represents:

$$\text{Unoptimized Bytes} - \text{Optimized Bytes}$$


- ❑ **Percent Savings:** A percentage value of bytes saved, calculated as:

$$\{ (\text{Client Bytes} - \text{Server Bytes}) / \text{Client Bytes} \} * 100$$

In the Savings panel shown above, the **Percent Savings** for FTP is **50%** and bandwidth savings is **2x**, which is calculated as *Client Bytes*/*Server Bytes*.

Note: The graph in the percent savings column represents savings over the last hour, while the label reflects the percent savings in the last minute. For more information on bandwidth savings, click on any row and navigate to the **Statistics > Traffic History** page. By default, the traffic history page displays bandwidth usage and bandwidth gain statistics for the corresponding service over the last hour.

The Interface Utilization panel displays statistics on interface use, reveals network performance issues, if any, and helps determine the need to expand your network.




- ❑ **Interface:** The interfaces are labeled with an adapter number followed by an interface number. For example, on 2-port bridge cards, the interface number is 0 for WAN and 1 for LAN connections; 4-port bridge cards have 0 and 2 for WAN and 1 and 3 for LAN.
- ❑ **Link state:** Indicates whether the interface is in use and functioning. It also displays the duplex settings and includes the following information:
 - **Up or Down:** **Up** indicates that the link is enabled and can receive and transmit traffic. **Down** indicates that the link is disabled and cannot pass traffic.
 - **Auto or Manual:** Indicates whether the link is auto-negotiated or manually set
 - **10Mbps, 100 Mbps or 1Gbps:** Displays the capacity of the link.
 - **FDX or HDX:** Indicates whether the interface uses full duplex or half duplex connection, respectively. In some cases, if a duplex mismatch occurs when the interface is auto-negotiated and the connection is set to half-duplex, the display icon changes to a yellow warning triangle. If you view a duplex mismatch, you can adjust the interface settings on the ProxySG in the **Configuration > Network > Adapters** tab.
- ❑ **Transmit Rate and Receive Rate:** Displays number of bits processed per second, on each interface.
The graphs in the transmit rate and receive rate columns represent interface activity over the last hour, while the value in the label represents interface activity over the last minute.
- ❑ **Errors:** Displays the number of transmission errors, if any, in the last hour. Interfaces with input or output errors are displayed in red.

For more information on an interface, click on any row; the **Statistics > Network > Interface History** page displays.


Monitoring System Resources and Connectivity Metrics

The **Statistics > Summary > Device** tab displays a snapshot of the key system resources, identification specifics, and the status of external devices that are connected to the ProxySG.

The identification panel provides information on the name of the ProxySG, IP address, hardware serial number, software version and the build (release) ID. You can copy and paste the information on this panel, into an email for example, when communicating with Symantec Support.



This information is also displayed on the Management Console banner and under **Configuration > General > Identification**. To assign a name to your ProxySG, see ["Configuring the ProxySG Name"](#) on page 35.



The **Statistics** area displays the current percentages of CPU usage and memory utilization, and the number of concurrent users. Concurrent users represents the number of unique IP addresses that are being intercepted by the ProxySG. For more information on these key resources, click the link; the corresponding panel under **Statistics > System > Resources** displays.

The Statistics panel also displays whether the ProxySG is enabled to:


- ☐ participate in an Application Delivery Network (ADN)
- ☐ serve as a ProxyClient Manager

The status information displayed for ADN and ProxyClient include the following options:

Feature	Status	Description
ADN	<i>Disabled</i>	<i>This ProxySG is not participating in an Application Delivery Network.</i>
	<i>Open ADN</i>	<i>This ProxySG is an ADN peer and can form a tunnel connection with any other ADN peer. An ADN Manager is not required for Open ADN.</i>
	<i>Configured as a Manager</i>	<i>This ProxySG serves as an ADN Manager.</i>

	<i>Connected to Managers</i>	<i>ADN is enabled and this ProxySG is connected to the Primary and the Backup ADN Manager.</i>
	<i>Connected to Primary Manager</i>	<i>ADN is enabled and this ProxySG is connected to the Primary ADN Manager.</i>
	<i>Connected to Backup Manager</i>	<i>ADN is enabled and this ProxySG is connected to the Backup ADN Manager. Implication: This ProxySG is unable to connect to the Primary ADN Manager. Inspect the Primary ADN Manager configuration in the Configuration > ADN > General tab.</i>
	<i>Not Connected to Either Manager</i>	<i>Although ADN is enabled, this ProxySG is not connected to the Primary or the Backup ADN Manager. Implication: The ADN is not functioning properly. Inspect the Primary and the Backup ADN Manager configuration in the Configuration > ADN > General tab.</i>
<i>ProxyClient</i>	<i>Client Manager Enabled; <number> Active Clients</i>	<i>This ProxySG serves as a ProxyClient Manager. Also displayed is the number of active clients that are connected to this ProxyClient Manager.</i>
	<i>Disabled</i>	<i>This ProxySG is not configured as a ProxyClient Manager.</i>

The **Connectivity** area displays the status of external devices and services that the ProxySG relies on, for an effective performance. The status indicates whether the ProxySG is able to communicate with the external devices and services that are configured on it.



External Devices	Status
WCCP	2 of 2 WCCP service groups ready
ICAP	4 of 4 ICAP health checks reporting Critical
DNS	1 of 1 DNS server reporting OK

The external devices or services, that can be configured on the ProxySG, include:

- ❑ WCCP capable routers/switches
- ❑ External ICAP devices (such as Blue Coat ProxyAV or Content Analysis appliances)
- ❑ DNS Servers
- ❑ Authentication realms

Only those external devices or services that are configured on the ProxySG are displayed on this panel. If, for example, ICAP is not yet enabled on the ProxySG, ICAP is not listed in the connectivity panel.

The connectivity status for these external devices is represented with an icon — Ok, Warning, or Critical. The icon and the text portray the most severe health status, after considering all the health checks configured, for the device or service.

With the exception of WCCP, click on any row to view the health status details in the **Statistics > Health Checks** tab. The **Statistics > Health Checks** tab provides information on the general health of the external services configured on the ProxySG, allows you to perform routine maintenance tasks and to diagnose potential problems. For more information on health checks, see ["Verifying the Health of Services Configured on the ProxySG"](#) on page 1353.


To view details on the status of WCCP capable devices in your network, click on the WCCP service row, the **Statistics > Network > WCCP** tab displays. The **Statistics > Network > WCCP** tab provides information on the configured service groups and their operational status. For more information on how to configure WCCP on the ProxySG, see [Chapter 33: "WCCP Configuration"](#) on page 775. For more detailed information about WCCP, refer to the *WCCP Reference Guide*.

Logging Out of the Management Console

To exit the current session, click the **Log Out** link on the Management Console banner.

You may be logged out of the ProxySG automatically, when a session timeout occurs. This security feature logs the user out when the Management Console is not actively being used. For more information, see ["Changing the ProxySG Timeout"](#) on page 37.

Thirty seconds before the session times out, the console displays a warning dialog. Click the **Keep Working** button or the **X** in the upper-right corner of the dialog box to keep the session alive.



If you do not respond within the thirty-second period, you are logged out and lose all the changes since the last submittal.

To log in again, click the hyperlink (**You need to log in again to use the console**) in the browser.

Note: If you do not want to log in again, close the browser window (not just the browser tab) to log out completely.

Accessing the ProxySG Using the CLI

You can connect to the ProxySG command-line interface via Secure Shell (SSH) using the IP address, username, password that you defined during initial configuration. The SSH management console service is configured and enabled to use SSHv2 and a default SSH host key by default. If you wish to access the CLI, you can use SSHv2 to connect to the ProxySG. An SSH host key for SSHv2 and an SSH management service are configured by default. If you want to use SSHv1 or Telnet without additional configuration.

Note: You can also access the CLI using Telnet or SSH v1. However, these management services are not configured by default. For instructions on configuring management services, see [Chapter 69: "Configuring Management Services"](#) on page 1267.

To log in to the CLI, you must have:

- ❑ the account name that has been established on the ProxySG
- ❑ the IP address of the ProxySG
- ❑ the port number (22 is the default port number)

SGOS supports different levels of command security:

- ❑ Standard, or unprivileged, mode is read-only. You can see but not change system settings and configurations. This is the level you enter when you first access the CLI.
- ❑ Enabled, or privileged, mode is read-write. You can make immediate but not permanent changes to the ProxySG, such as restarting the system. This is the level you enter when you first access the Management Console.
- ❑ Configuration mode allows you to make permanent changes to the ProxySG configuration. To access Configuration mode, you must be in Enabled mode.

When you log in to the Management Console using your username and password, you are directly in configuration mode.

However, if you use the CLI, you must enter each level separately:

```
Username: admin
Password:
SGOS> enable
Enable Password:
SGOS# configure terminal
Enter configuration commands, one per line. End with CTRL-Z.
SGOS#(config)
```

For detailed information about the CLI and the CLI commands, refer to the *Command Line Interface Reference*.

Note: Most tasks can be performed in both the Management Console and the CLI. This guide covers procedures for the Management Console; refer to the *Command Line Interface Reference* for related CLI tasks. Tasks that are available only in the Management Console or only in the CLI are noted as such.

Section A: Configuring Basic Settings

This sections describes how to configure basic settings, such as the ProxySG name time settings, and login parameters. It includes the following topics:

- ❑ ["How Do I...?"](#) on page 34
- ❑ ["Configuring the ProxySG Name"](#) on page 35
- ❑ ["Changing the Login Parameters"](#) on page 35
- ❑ ["Viewing the Appliance Serial Number"](#) on page 38
- ❑ ["Configuring the System Time"](#) on page 38
- ❑ ["Synchronizing to the Network Time Protocol"](#) on page 40

How Do I...?

To navigate this section, identify the task to perform and click the link:

How do I...?	See...
Assign a name to identify the ProxySG?	"Configuring the ProxySG Name" on page 35
Change the logon parameters?	"Changing the Login Parameters" on page 35
Locate the Appliance Serial Number?	"Viewing the Appliance Serial Number" on page 38
Configure the local time on the ProxySG?	"Configuring the System Time" on page 38
Synchronize the ProxySG to use the Network Time Protocol (NTP)?	"Synchronizing to the Network Time Protocol" on page 40
Change the log-in username and password?	"Changing the Administrator Account Credentials" on page 35
Configure a console realm name to identify the ProxySG that I am accessing (before I log in to the Management Console)?	"Changing the ProxySG Realm Name" on page 37
Configure the time for console log out on the ProxySG?	"Changing the ProxySG Timeout" on page 37

Configuring the ProxySG Name

You can assign any name to a ProxySG. A descriptive name helps identify the system.

To set the ProxySG name:

1. Select **Configuration > General > Identification**.
2. In the **Appliance name** field, enter a unique name for the appliance.
3. Click **Apply**.

Changing the Login Parameters

You can change the console username and password, the console realm name which displays when you log in to the ProxySG, and the auto-logout time on the ProxySG. The default value is 900 seconds.

The Management Console requires a valid administrator username and password to have full read-write access; you do not need to enter a privileged-mode password as you do when using the CLI. A privileged-mode password, however, must already be set.

Note: To prevent unauthorized access to the ProxySG, only give the console username and password to those who administer the system. You can also configure the appliance to lock out administrative users in the local realm if they exceed a number of failed authentication attempts. For more information, see ["Enhancing Security Settings for the Local User List"](#) on page 941.


Changing the Administrator Account Credentials

During the initial configuration of your ProxySG appliance, a console administrator username and password was created. This is a special account that can always be used to administer the appliance from either the web-based Management Console or the Command Line Interface. You can change the username and the password of this administrator account.

Note: Changing the console account's username or password causes the Management Console to refresh, requiring you to log in again using the new credentials. Each parameter must be changed and individually refreshed. You cannot change both parameters at the same time.

To change the username:

1. Select **Configuration > Authentication > Console Access > Console Account**.



2. Edit the username of the administrator that is authorized to view and revise console properties. Only one console account exists on the ProxySG. If you change the console account username, that username overwrites the existing console account username. The console account username can be changed to anything that is not null and contains no more than 64 characters.
3. Click **Apply**. After clicking **Apply**, an **Unable to Update configuration** error is displayed. This is expected: although the username change was successfully applied, the configuration could not be fetched from the ProxySG appliance because the old username was offered in the fetch request.
4. Refresh the screen. You are challenged for the new username.

To change the password:

The console password and privileged-mode password were defined during initial configuration of the system. The console password can be changed at any time. The privileged-mode, or enabled-mode, password can only be changed through the CLI or the serial console.

1. Select **Configuration > Authentication > Console Access > Console Account**.
2. Click **Change Password**.
3. Enter and re-enter the console password that is used to view and edit configuration information. The password must be from 1 to 64 characters long. As you enter the new password, it is obscured with asterisks. Click **OK**.

Note: This does not change the enabled-mode password. You can only change the enabled-mode password through the CLI.

4. Refresh the screen, which forces the SGOS software to re-evaluate current settings. When challenged, enter the new password.
5. (Optional) Restrict access by creating an access control list or by creating a policy file containing `<Admin>` layer rules. For more information, see ["Limiting Access to the ProxySG"](#) on page 59.

Changing the ProxySG Realm Name

When you have multiple ProxySG appliances in your network, you can configure a console realm name to identify the appliance that you are accessing.


When you log in to the Management Console, using a browser, the browser's pop-up dialog displays. This dialog identifies the ProxySG that is requesting the username and password.

If configured, the realm name displays on the pop-up dialog. The default realm name is usually the IP address of the ProxySG. You can, however, change the display string to reflect your description of the ProxySG.

To change the realm name:

1. Select **Configuration > Authentication > Console Access > Console Account**.
2. Enter a new realm name in **Console realm name**.
3. Click **Apply**.

The next time you log in to the Management Console, the new realm name displays on the browser's pop-up dialog.



Changing the ProxySG Timeout

The timeout is the length of time a Web or CLI session persists before you are logged out. The default timeout for these options is as follows:

- ☐ **Enforce Web auto-logout**—15 minutes
- ☐ **Enforce CLI auto-logout**—5 minutes

To change the timeout:

1. Select **Configuration > Authentication > Console Access > Console Account**.
2. Configure the timeout by doing one of the following:
 - Set values for the Web or CLI auto-logout. Acceptable values are between **1** and **1440** minutes.
 - Deselect the auto-timeout to disable it.
3. Click **Apply**.

Viewing the Appliance Serial Number

The ProxySG serial number assists Blue Coat Systems Customer Support when analyzing configuration information, including heartbeat reports. The appliance serial number is visible on the Management Console banner.

Configuring the System Time


To manage objects, the ProxySG must know the current Coordinated Universal Time (UTC), which is the international time standard and is based on a 24-hour clock. The ProxySG accesses the Network Time Protocol (NTP) servers to obtain accurate UTC time and synchronizes its time clock.

By default, the ProxySG connects to an NTP server in the order they are listed on the **NTP** tab and acquires the UTC time. You can view UTC time under **UTC** in the **Configuration > General > Clock > Clock tab**. If the appliance cannot access any of the listed NTP servers, you must manually set the UTC time.

You can, however, also record time stamps in local time. To record time stamps in local time, you must set the local time based on your time zone. The ProxySG appliance ships with a limited list of time zones. If a specific time zone is missing from the included list, you can update the list at your discretion. The list can be updated by downloading the full time zone database from <http://download.bluecoat.com/release/timezones.tar>. Also, the time zone database might need to be updated if the Daylight Savings rules change in your area.


To set local time:

1. Select **Configuration > General > Clock > Clock**.



The screenshot shows the 'Clock' configuration page. It has two tabs: 'Clock' and 'NTP'. The 'Clock' tab is active. Under 'Current time:', there are fields for 'UTC:', 'Local:', and 'Timezone:'. The 'UTC:' field shows '21:43:40' and '19 Jun 2008'. The 'Local:' field shows '21:43:40' and '19 Jun 2008'. The 'Timezone:' field shows 'UTC'. There is a 'Sync' button next to the UTC field and a 'Set Time zone' button next to the Timezone field. Below this, there is a section for 'Update time zone database' with an 'Installation URL:' field showing 'http://', an 'Install' button, and a 'Set to default' button. At the bottom, there is a section for 'Method for acquiring UTC:' with a checkbox for 'Enable NTP' (checked), a 'Query interval (minutes)' field showing '60', and an 'Acquire UTC time' button.

2. Click **Set Time zone**. The Time Zone Selection dialog displays.



3. Select the time zone that represents your local time. After you select the local time zone, event logs record the local time instead of GMT. To add additional time zones to the list, update the appliance's time zone database, as described in the following procedure.

4. Click **OK** to close the dialog.
5. Click **Apply**.

To update the database:

1. Select **Configuration > General > Clock > Clock**.
2. Enter the URL from which the database will be downloaded or click **Set to default**.
3. Click **Install**.

To acquire the UTC:

1. Ensure that **Enable NTP** is selected.
2. Click **Acquire UTC Time**.

Synchronizing to the Network Time Protocol

The Network Time Protocol (NTP) is used to synchronize the time of a computer client or server to another server or reference time source, such as a radio or satellite receiver or modem. There are more than 230 primary time servers, synchronized by radio, satellite and modem.


The ProxySG ships with a list of NTP servers available on the Internet, and attempts to connect to them in the order they appear in the NTP server list on the **NTP** tab. You can add others, delete NTP servers, and reorder the NTP server list to give a specific NTP server priority over others.

The ProxySG uses NTP and the Coordinated Universal Time (UTC) to keep the system time accurate.

You can add and reorder the list of NTP servers the ProxySG uses for acquiring the time. (The reorder feature is not available through the CLI.)

To add an NTP server:

1. Select **Configuration > General > Clock > NTP**.



2. Click **New**. The Add List Item dialog displays.
3. Choose one of the following:
 - **Domain name:** Enter a domain name of an NTP server that resolves to an IPv4 or IPv6 address.
 - **IP address:** Enter an IPv4 or IPv6 address of an NTP server.
4. Click **OK** to close the dialog.
5. Click **Apply**.

To change the access order:

NTP servers are accessed in the order displayed. You can organize the list of servers so the preferred server appears at the top of the list. This feature is not available through the CLI.

1. Select **Configuration > General > Clock > NTP**.
2. Select an NTP server to promote or demote.
3. Click **Promote entry** or **Demote entry** as appropriate.
4. Click **Apply**.

Appendix: Required Ports, Protocols, and Services

Depending on your ProxySG appliance configuration, you must open certain ports and protocols on your firewalls for the appliance to function as intended, or to allow connectivity to various components and data centers. For full details, refer to the following knowledge base article:

<https://www.symantec.com/docs/INFO5294>

Chapter 3: Licensing

This section describes the ProxySG licensing behavior and includes the following topics:

- ❑ "About Licensing"
- ❑ "Disabling the Components Running in Trial Period" on page 48
- ❑ "Registering and Licensing the Appliance" on page 49
- ❑ "Enabling Automatic License Updates" on page 56
- ❑ "Viewing the Current License Status" on page 57

About Licensing

Each ProxySG appliance requires a license to function. The license is associated with an individual ProxySG serial number and determines what software features are available and the number of concurrent users that are supported.

When you configure a new hardware appliance, the Blue Coat ProxySG configuration wizard automatically installs a trial license that allows you to use all software features with support for an unlimited number of concurrent users for 60 days. (Trial periods are not applicable to ProxySG virtual or Advanced Secure Gateway appliances.) The software features that are available depend on what license edition is installed and what license features you have purchased.

The following sections describe the licensing options:

- ❑ "License Expiration" on page 48
- ❑ "Licence Types" on page 46
- ❑ "License Expiration" on page 48

Note: The information in this chapter does not apply to the Secure Web Gateway Virtual Appliance (SWG VA). For licensing and upgrade information specific to the SWG VA, refer to the *Secure Web Gateway Initial Configuration Guide*.

License Editions

The license edition determines what features are available. The ProxySG supports two license editions:

- ❑ **Proxy Edition License**—Supports all security and acceleration features. The Proxy Edition allows you to secure Web communications and accelerate the delivery of business applications.

- ❑ **MACH5 Edition License**—Supports acceleration features and Symantec Cloud Service; on-box security features are not included in this edition. The MACH5 base license allows acceleration of HTTP, FTP, CIFS, DNS, MAPI, and streaming protocols.

During the setup process, you indicate how you will deploy the appliance, which determines trial license edition is installed. If you indicate that you will be using the appliance as an acceleration node, a MACH5 trial license is installed. For other deployment types, the wizard prompts you to select Proxy edition.

Proxy Edition and MACH5 license edition can run on any ProxySG platform. The only differences are the supported software features and the default configuration settings. These differences are described in the following sections:

- ❑ ["Differences in Default Configuration Settings"](#)
- ❑ ["MACH5 Feature Set"](#) on page 45
- ❑ ["Switching Between the License Editions"](#) on page 46

Differences in Default Configuration Settings

Because the different license editions are intended for different deployments, some of the default configuration settings are different between license editions. The Proxy Edition is meant to provide security and is thus more restrictive in allowing traffic through whereas the MACH5 edition is geared for application acceleration and is therefore more permissive. The difference in the defaults are as follows:

- ❑ Default policy on the ProxySG: This setting determines whether, by default, all traffic is allowed access or denied access to requested content.
 - MACH5 Edition: Allow
 - Proxy Edition: Deny
- ❑ Trust destination IP provided by the client: (only applicable for transparent proxy deployments) This setting determines whether or not the ProxySG will perform a DNS lookup for the destination IP address that the client provides.
 - MACH5 Edition: Enabled. The proxy trusts the destination IP included in the client request and forwards the request to the OCS or services it from cache.
 - Proxy Edition: Disabled
- ❑ HTTP tolerant request parsing: The tolerant HTTP request parsing flag causes certain types of malformed requests to be processed instead of being rejected.
 - MACH5 Edition: Enabled. Malformed HTTP requests are not blocked.
 - Proxy Edition: Disabled
- ❑ Transparent WAN intercept on bridge cards: This setting indicates whether the proxy should intercept or bypass packets on the WAN interface.
 - MACH5 Edition: Bypass transparent interception
 - Proxy Edition: Allow transparent interception

- ❑ Resource overflow action: This setting indicates whether the proxy should bypass or drop new connections when resources are scarce.
 - MACH5 Edition: Bypass
 - Proxy Edition: Drop

MACH5 Feature Set

The MACH5 license edition provides a subset of the full feature set provided by the Proxy Edition license. The following table describes feature support on an appliance running a MACH5 license:

Table 3–1 MACH5 Feature Support

Feature	MACH5 Support
Access Logging	Supported; CIFS, Endpoint Mapper, FTP, HTTP, TCP Tunnel, Windows Media, Real Media/QuickTime, SSL, HTTPS Forward Proxy, MAPI and Flash
ADN	Supported
Authentication	On-box authentication supported for administrative access (IWA, LDAP, RADIUS, SiteMinder, COREid, and local realms only). User authentication is not supported on-box except when combined with Symantec Cloud Service. When using the Web Security Module of the Symantec Cloud Service, LDAP and IWA are supported to provide user authentication details for cloud-based policy enforcement.
Bandwidth Management	Supported
Content Filtering	Not supported on-box; Use Symantec Cloud Security Services for Content Filtering.
External Services(ICAP)	Not supported
Forwarding	Forwarding hosts: Supported SOCKS: Not supported
HTTP Compression	Supported
Instant Messaging	Not supported
Peer-to-Peer	Not supported
Policy Controls	Acceleration-based policy controls: Supported Exception pages: Not supported
ProxyClient	Acceleration: Supported Content Filtering: Not Supported

Table 3–1 MACH5 Feature Support (Continued)

Feature	MACH5 Support
Proxy Services	CIFS, FTP, HTTP, MAPI and Streaming (Windows Media, Real Media and QuickTime) are Supported. Flash proxy is also supported, however you must purchase and install an add-on license to use this service. SSL Termination is also supported. Some appliance models include an SSL license; other models require that you purchase and install an add-on license.
Threat Protection Services	Not supported

Switching Between the License Editions

This section describes the effects of switching between the license editions.

- ❑ **Upgrading from the MACH5 Edition to the Proxy Edition**—You can upgrade from the MACH5 Edition license to the Proxy Edition license at any time, as long as you use the same hardware. Upon upgrade, the entire license file is regenerated. This is because the defaults must be readjusted to reflect the change in functionality, and must include some proxy-specific configurations, such as advanced services and access logging logs and formats, which are added during the upgrade.

Note: The existing configuration is not changed during the upgrade.

All the MACH5 Edition functionality is supported in the Proxy Edition, so an upgrade does not affect CLI or policy commands.

- ❑ **Downgrading from a Proxy Edition to a MACH5 Edition**—You must install a new license to switch from a Proxy Edition license to a MACH5 Edition license. This license downgrade can be performed only by restoring the ProxySG appliance to its factory defaults; as a result, your existing configuration will be deleted and you will have to reconfigure the appliance.

Licence Types

There are several different types of licenses:

- ❑ **Trial**—The 60-day license that ships with new ProxySG physical appliances. (Trial licenses are not available on virtual or Advanced Secure Gateway appliances.) All licensable components for the trial edition are active and available to use. In addition, the Base SGOS user limit is unlimited. When a full license is installed, any user limits imposed by that license are enforced, even if the trial period is still valid.
- ❑ **Demo**—A temporary license that can be requested from Symantec to extend the evaluation period.

- ❑ **Permanent**—A license for hardware platforms that permanently unlocks the software features you have purchased. When a permanent license is installed, any user limits imposed by that license are enforced, even if the trial period is still valid.
- ❑ **Subscription-based**—A license that is valid for a set period of time. After you have installed the license, the ProxySG will have full functionality, and you will have access to software upgrades and product support for the subscription period.

Note: When a full license (permanent or subscription-based) or demo license is installed during the trial period, components previously available in the trial period, but not part of that license, remain available and active for the remainder of the trial period. However, if the license edition is different than the trial edition you selected, only functionality available in the edition specified in the license remains available for trial. If you do not want the trial components to be available after you install a full license, you can disable them. See ["Disabling the Components Running in Trial Period"](#) on page 48 for instructions.

Licensing Terms

ProxySG Appliances

Within sixty (60) days of the date from which the user powers up the ProxySG appliance ("Activation Period"), the Administrator must complete the ProxySG licensing requirements as instructed by the ProxySG to continue to use all of the ProxySG features. Prior to the expiration of the Activation Period, the ProxySG software will deliver notices to install the license each time the Administrator logs in to manage the product. Failure to install the license prior to the expiration of the Activation Period may result in some ProxySG features becoming inoperable until the Administrator has completed licensing.

Proxy Client

The Administrator may install the Proxy Client only on the number of personal computers licensed to them. Each personal computer shall count as one "user" or "seat." The ProxyClient software may only be used with Blue Coat ProxySG appliances. The Administrator shall require each user of the Symantec ProxyClient software to agree to a license agreement that is at least as protective of Symantec and the Symantec ProxyClient software as the Symantec EULA.

ProxySG Virtual Appliances, MACH5 or Secure Web Gateway (SWG) Edition:

The ProxySG Virtual Appliances (MACH5 or Secure Web Gateway edition) are licensed on either a perpetual or subscription basis for a maximum number of concurrent users. Support for the Virtual Appliances will be subject to the separate support agreement entered into by the parties if the Administrator licenses the Virtual Appliances on a perpetual basis. The Virtual Appliances will (a) not function upon expiration of the subscription if the Administrator licenses the Virtual Appliances on a subscription basis; or (b) if the traffic exceeds the

maximum number of concurrent users/connections, features may not function beyond the maximum number of concurrent users/connections. This means that, in these cases, the network traffic will only be affected by the default policy set by the Administrator (either pass or deny). Such cessation of functionality is by design, and is not a defect in the Virtual Appliances. The Administrator may not install the same license key or serial number on more than one instance of the Virtual Appliance. The Administrator may move the Virtual Appliance along with its license key and serial number to a different server, provided that server is also owned by the Administrator and the Administrator permanently deletes the prior instance of the Virtual Appliance on the server on which it was prior installed. The Virtual Appliances require a third party environment that includes software and/or hardware not provided by Symantec, which the Administrator will purchase or license separately. Symantec has no liability for such third party products.

License Expiration

When any licensed component expires, those components do not process requests; all requests bypass the ProxySG if the default policy is set to **Allow**. A license expiration notification message is logged in the Event Log (see "[Viewing Event Log Configuration and Content](#)" on page 1316 for details on how to view the event log).

If a license expires, users might not receive notification, depending upon the application they are using. Notifications do occur for the following:

- ❑ HTTP (Web browsers)—An HTML page is displayed stating the license has expired.
- ❑ SSL—An exception page appears when an HTTPS connection is attempted, but only if the appliance is deployed explicitly or in the case of transparent proxy deployments, SSL interception is configured.
- ❑ FTP clients—If the FTP client supports it, a message is displayed stating the license has expired.
- ❑ Streaming media clients—If the Windows Media Player, RealPlayer, or QuickTime player version supports it, a message is displayed stating the license has expired.
- ❑ ProxyClient—After the trial license has expired, clients cannot connect to the ADN network.
- ❑ You can still perform configuration tasks through the CLI, SSH console, serial console, or Telnet connection. Although the component is disabled, feature configurations are *not* altered. Also, policy restrictions remain independent of component availability.

Disabling the Components Running in Trial Period

You have the option to disable access to features that are running in trial period; however, you cannot selectively disable trial period features. You must either enable all of them or disable all of them.

Note: Because licensing trial periods are not offered on the ProxySG VA, this option is not available on virtual appliances.

To disable trial period components:

1. Select **Maintenance > Licensing > View**.
2. Select the **Trial Components are enabled** option.
3. Click **Apply**.
4. Click **Refresh Data**. All licenses that are in trial period switch from **Yes** to **No**. Users cannot use these features, and no dialogs warning of license expiration are sent.

Also notice that this option text changes to **Trial Components are disabled: Enabled**. Repeat this process to re-enable trial licenses.

Registering and Licensing the Appliance

Before you can register and license your appliance, you must have the following:

- ❑ The serial number of your appliance. See "[Locating the System Serial Number](#)" on page 49.
- ❑ A BlueTouch Online account. See "[Obtaining a BlueTouch Online Account](#)" on page 50.

You can then register the appliance and install the license key. The following sections describe the available options for completing the licensing process:

- ❑ If you have not manually registered the appliance, you can automatically register the appliance and install the software license in one step. See "[Registering and Licensing Blue Coat Appliance and Software](#)" on page 50.
- ❑ If you have a new appliance that previously has been registered, the license is already associated with the appliance. In this case you just need to retrieve the license. See "[Installing a License on a Registered System](#)" on page 51.
- ❑ If you have older hardware that previously has been registered or if the ProxySG does not have Internet access, you must install the license manually. See "[Manually Installing the License](#)" on page 52.
- ❑ After the initial license installation, you might decide to use another feature that requires a license. The license must be updated to support the new feature.

Locating the System Serial Number

Each ProxySG serial number is the appliance identifier used to assign a license key file. The ProxySG contains an EEPROM with the serial number encoded. The appliance recognizes the serial number upon system boot-up. The appliance serial number is located in the information bar at the top of the Management Console.

Serial numbers are not pre-assigned on the ProxySG Virtual Appliance. You retrieve the serial number from the Symantec Licensing Portal, and enter the serial number during initial configuration. Refer to the *ProxySG VA Initial Configuration Guide* for more information.

Obtaining a BlueTouch Online Account

Before you can register your ProxySG and retrieve the license key, you must have a Blue Coat BlueTouch Online user account.

If you do not have a BlueTouch Online account or have forgotten your account information, perform the following procedure.

To obtain a BlueTouch Online account:

1. Select the **Maintenance > Licensing > Install** tab.
2. In the **License Administration** field, click **Register/Manage**. The License Configuration and Management System Web page displays.
3. Perform one of the following:
 - To obtain a new account, click the link for **Need a BlueTouch Online User ID. Under Login Assistance (BlueTouch Online)**, click **Request Login User ID/Password**. Fill out the Web form; BlueTouch Online information will be sent to you.
 - To obtain your current information for an existing account, click the **Forgot your password** link.

Registering and Licensing Blue Coat Appliance and Software

If you have not manually registered the appliance, you can automatically register the appliance and install the software license in one step as described in the following procedure.

To register the appliance and software:

1. In a browser, go to the following URL to launch the Management Console:
`https://appliance_IP_Address:8082`
2. Enter the access credentials specified during initial setup.
3. Click **Management Console**. The **License Warning** tab displays.

License Warning

This device is operating in the trial period. Trial expiration date is 2007-11-10

Hardware Registration

☒ Register hardware with Blue Coat automatically

WebPower User ID: JENNIFER.CUSTOMER

Password: *****

☐ Hardware has been manually registered

Registration Status: Hardware auto-registration and license retrieval in progress. The license retrieval may take up to a minute.

4. Make sure the **Register hardware with Blue Coat automatically** radio button is selected.
5. Enter your BlueTouch Online credentials and click **Register Now**. This opens a new browser page where you complete the registration process. When the hardware is successfully registered, the **Registration Status** field on the License Warning tab will display the Hardware auto-registration successful message. You can close the new browser tab or window that displays the License Self-service page.
6. Click **Continue**.

Installing a License on a Registered System

If the ProxySG is a new system and the appliance has been registered, retrieve the associated license by completing this procedure.

To retrieve the software license:

1. Select the **Maintenance > Licensing > Install** tab.
2. Click **Retrieve**. The Request License Key dialog displays.

Request License Key

BlueTouch Online

Sign in to your BlueTouch Online account to retrieve your license key

User ID: hub.porter@company.com

Password: *****

Do you need a BlueTouch Online User ID?
Did you forget your password?

3a

3b

Request license Show results Close

3. Enter information:
 - a. Enter your BlueTouch Online account login information.
 - b. Click **Request License**. The Confirm License Install dialog box displays.
 - c. Click **OK** to begin license retrieval (the dialog closes).
4. (Optional) Click **Show results** to verify a successful retrieval. If any errors occur, check the ability for the ProxySG to connect to Internet.
5. Click **Close** to close the Request License Key dialog.
6. To validate the license, restart the appliance.
In the Management Console, select **Maintenance > Tasks**.
 - Click **Hardware and Software**.
 - Click **Restart now**.

Manually Installing the License

Perform manual license installation if:

- ☐ The ProxySG serial number is not associated with a software license (you have registered the hardware separately)
- ☐ The ProxySG does not have Internet access

Note: Locate the email from Symantec that contains the activation code(s) for your software. You require these activation codes, as well as your appliance serial number, to complete the licensing process on the Symantec Licensing Portal.

Manually retrieve and install the license:

Tip: Follow these steps if you registered the hardware separately.


1. In the Management Console, select **Maintenance > Licensing > Install**.
2. Click **Register/Manage**. The licensing portal opens in a browser window and prompts you for your BlueTouch Online login information.
3. Enter your login credentials and click **Login**. The Licensing Portal prompts you to enter your activation code.
4. Enter the activation code and follow the prompts to complete the process.
When prompted to accept the license agreement, read and accept the terms.
The software license is now associated with the appliance.
5. (If necessary) Repeat the previous steps for your other activation codes.
6. To validate the license, restart the appliance.
In the Management Console, select **Maintenance > Tasks**.
 - Click **Hardware and Software**.
 - Click **Restart now**.

Download and manually install the license:

Tip: Follow these steps if the appliance does not have access to the Internet.


1. In the activation email, click the link to the Licensing Portal. The browser opens the portal on the main page.
2. Select **ProxySG > License Download**. The portal prompts you for your appliance serial number.

BLUE COAT



The screenshot displays the Blue Coat Licensing Portal. At the top is a dark blue navigation bar with links: SOLUTIONS, PRODUCTS, CUSTOMERS, RESOURCES, PARTNERS, SUPPORT, NEWS, and COMPANY. Below this is a header section with a small image on the left and the text "Blue Coat Licensing Portal" on the right, with a "Logout" link. A left-hand menu lists various options under the "REGISTRATION" heading, including Home, CacheFlow, Content Analysis System, Content Filtering, Director, IntelligenceCenter / PolicyCenter, Management Center, On Premise Inspection Connector, PacketShaper, PacketShaper S-Series, ProxyAV, ProxySG, Software Add-On Activation, ProxySG MACH5 Edition VA / Secure Web Gateway VA, Upgrades, Cold Standby Swap, License Download, Generate Trial License, Reporter, SSL Visibility, and Logout. The main content area is titled "ProxySG License Download" and contains the instruction "Please enter your hardware serial number". Below this is a red label "Serial Number" followed by a text input field. At the bottom of the form are two buttons: "Clear form" and "Submit".

- Follow the prompts to enter your serial number and download the license file.



BLUE COAT

SOLUTIONS PRODUCTS CUSTOMERS RESOURCES PARTNERS SUPPORT NEWS COMPANY

Blue Coat Licensing Portal [Logout](#)

REGISTRATION

- Home
- CacheFlow
- Content Analysis System
- Content Filtering
- Director
- IntelligenceCenter / PolicyCenter
- Management Center
- On Premise Inspection Connector
- PacketShaper
- PacketShaper S-Series
- ProxyAV
- ProxySG
 - Software Add-On Activation
 - ProxySG MACH5 Edition VA / Secure Web Gateway VA
 - Upgrades
 - Cold Standby Swap
 - License Download
 - Generate Trial License
- Reporter
- SSL Visibility
- Logout

ProxySG License Download

You can obtain your license file below

Serial Number: 1012162236
 Product Description: Blue Coat SG300-25, Proxy Edition
 Download: [Download License File](#)

Product Description	Expiration Date
CachePulse Optimization Service, SG300 MACH5/Proxy Edition, 1 yr subscription	2015-10-27
Flash Proxy License, SG300	2019-12-31
Software, SGOS 6.x, Model SG300-25-PR	2019-12-31
Encrypted Tap license, SG300 Proxy Edition	2019-12-31
Web Application Protections, SG300 Proxy Edition, 1 yr subscription	2015-10-27

- Save the license file to a location that your appliance can access.
- In the Management Console, select **Maintenance > Licensing > Install**, and then select the appropriate option from the **License Key Manual Installation** drop-down list:

Note: A message is written to the event log when you install a license through the ProxySG.

- Remote URL**—Choose this option if the file resides on a Web server; then click **Continue**. The console displays the Install License Key dialog. Enter the URL path and click **Install**. When installation is complete, click **OK**.
- Local File**—Choose this option if the file resides in a local directory; then click **Continue**. The Open window displays. Navigate to the license file and click **Open**. When installation is complete, click **OK**.

6. To validate the license, restart the appliance.
In the Management Console, select **Maintenance > Tasks**.
 - Click **Hardware and Software**.
 - Click **Restart now**.

Adding an Add-on License

If you purchased a supplemental license to enable add-on features, you must update the ProxySG license by logging into the Symantec Licensing Portal and generating the license activation code. To do this, you must have the code for your ordered add-on feature that was sent in the e-mail from Symantec and the hardware serial number of the ProxySG that is to run the add-on feature.

To add a supplemental license:

1. Obtain the e-mail sent by Symantec that contains the license activation code(s) for the add-on license.
2. Log on to the Network Protection Licensing Portal:
https://services.bluecoat.com/eservice_enu/licensing/sso.cgi
3. Click **Licensing**.
4. Click the **License Your Blue Coat Products** link. The browser opens the licensing portal. If the portal prompts you to use your credentials again, enter them. The browser displays the portal home page.
5. In the **Enter Activation Code** field, enter the add-on product code from the e-mail; click **Next**. The Licensing Portal displays the Software Add-On Activation page.
6. In the Appliance Serial Number field, enter the ProxySG serial number. Click **Submit**.
7. The portal displays the license agreement; read and accept the agreement.

The portal displays a screen with license details for the software add-on. You can click **Back** and proceed to the next section.

BLUE COAT

The screenshot shows the Blue Coat Licensing Portal interface. At the top is a navigation bar with links: SOLUTIONS, PRODUCTS, CUSTOMERS, RESOURCES, PARTNERS, SUPPORT, NEWS, and COMPANY. Below this is a header section with the Blue Coat logo and the title 'Blue Coat Licensing Portal'. A 'Logout' link is in the top right corner. On the left is a sidebar menu with various options, including 'Software Add-On Activation' which is currently selected. The main content area is titled 'Software Add-On Activation' and displays the following information:

Here is your Blue Coat CPOS Proxy License:

License type:	CachePulse Optimization Service, SG300 MACH5/Proxy Edition, 1 yr subscription
License:	63362_2X2MT

Below the license information, there is a note: 'More information on license configuration can be found at the following URL: <http://www.bluecoat.com/support/manuals>'. At the bottom of the main content area is a 'Back' button.

Adding the Add-on License to the ProxySG

You must retrieve the updated license to the ProxySG.

To update the license:

1. From the ProxySG Management Console, select the **Maintenance > Licensing > Install** tab.
2. Click **Retrieve**. The appliance retrieves the license.
3. To verify a successful license update, select the **Licensing > View** tab; the console displays the new license in the General License Information section.

Enabling Automatic License Updates

The license automatic update feature allows the ProxySG to contact the Blue Coat licensing Web page 31 days before the license is to expire. If a new license has been purchased and authorized, the license is automatically downloaded. If a new license is not available on the Web site, the ProxySG continues to contact the Web site daily for a new license until the current license expires. Outside the above license expiration window, the ProxySG performs this connection once every 30 days to check for new license authorizations. This feature is enabled by default.

To configure the license auto-update:

1. Select the **Maintenance > Licensing > Install** tab.
2. Select **Use Auto-Update**.
3. Select **Apply**.
4. You must log in to your License Management account:

`https://services.bluecoat.com/eservice_enu/licensing/mgr.cgi`

5. Click **Update License Key**.

Viewing the Current License Status

You can view the license status in the Management Console in the following ways:

- ❑ Select **Statistics > Configuration > Maintenance**. The license status displays as a link in the upper right hand-corner. Hovering over the license link displays information, such as the expiration date of the trial period. Click the link to switch to the **View** license tab.
- ❑ Select **Maintenance > Licensing > View**. The tab displays the license components with expiration dates.
- ❑ Select **Maintenance > Health Monitoring**. The tab displays thresholds for license expiration dates.

The screenshot shows the 'View' tab of the License Management console. It has two main sections: 'General License Information' and 'Licensed Components'.

General License Information:

- Appliance serial number: 123456789
- License creation date: 2009-07-21
- Concurrent users (with ADN enabled): 50
- Concurrent users (without ADN enabled): 150

Licensed Components:

Component	Valid	Expiration Date
SGOS 5 MACH5 Edition	yes	2009-06-23
Windows Media Streaming	yes	2009-06-23
Real Media Streaming	yes	2009-06-23
QuickTime Streaming	yes	2009-06-23
SSL	yes	2009-06-23

Below the table are buttons for 'View Details' and 'Refresh Data'. At the bottom, there is a checkbox for 'Trial Components are enabled' which is currently unchecked.

Annotations:

- 'Current high-level license data' points to the 'General License Information' section.
- 'License components' points to the 'Licensed Components' table.
- 'For more details, select a component and click.' points to the 'View Details' button.

Each licensable component is listed, along with its validity and its expiration date.

- To view the most current information, click **Refresh Data**.
- Highlight a license component and click **View Details**. A dialog displays with more detailed information about that component.
- If the trial period is enabled and you click **Maintenance > Licensing > View**, the Management Console displays an option to disable the trial components. If the trial period is disabled, the Management Console displays an option to enable the trial components.

See Also

- ❑ ["About Licensing" on page 43](#)
- ❑ ["Disabling the Components Running in Trial Period" on page 48](#)
- ❑ ["Locating the System Serial Number" on page 49](#)
- ❑ ["Obtaining a BlueTouch Online Account" on page 50](#)
- ❑ ["Registering and Licensing Blue Coat Appliance and Software" on page 50](#)

Chapter 4: Controlling Access to the ProxySG

This section describes how to control user access to the ProxySG. It includes the following topics:

- ❑ ["Limiting Access to the ProxySG"](#) on page 59
- ❑ ["About Password Security"](#) on page 60
- ❑ ["Limiting User Access to the ProxySG—Overview"](#) on page 61
- ❑ ["Moderate Security: Restricting Management Console Access Through the Console Access Control List \(ACL\)"](#) on page 63
- ❑ ["Maximum Security: Administrative Authentication and Authorization Policy"](#) on page 64

Limiting Access to the ProxySG

You can limit access to the ProxySG by:

- ❑ Restricting physical access to the system and by requiring a PIN to access the front panel.
- ❑ Restricting the IP addresses that are permitted to connect to the ProxySG CLI.
- ❑ Requiring a password to secure the Setup Console.
- ❑ Configuring the appliance to lock out users in the local realm if they exceed a number of failed authentication attempts. For more information, see ["Enhancing Security Settings for the Local User List"](#) on page 941.

These safeguards are in addition to the restrictions placed on the console account (a console account user password) and the Enable password.

By using every possible method (physically limiting access, limiting workstation IP addresses, and using passwords), the ProxySG is very secure.

This section discusses:

- ❑ ["Requiring a PIN for the Front Panel"](#)
- ❑ ["Limiting Workstation Access"](#) on page 60
- ❑ ["Securing the Serial Port"](#) on page 60

Requiring a PIN for the Front Panel

On systems that have a front panel display, you can create a four-digit PIN to protect the system from unauthorized use. The PIN is hashed and stored. You can only create a PIN from the command line.

To create a front panel PIN, after initial configuration is complete:

From the (config) prompt:

```
SGOS#(config) security front-panel-pin PIN
```

where *PIN* is a four-digit number.

To clear the front-panel PIN, enter:

```
SGOS#(config) security front-panel-pin 0000
```

Limiting Workstation Access

During initial configuration, you have the option of preventing workstations with unauthorized IP addresses from accessing the ProxySG for administrative purposes. This covers all access methods - Telnet, SNMP, HTTP, HTTPS and SSH. If this option is not enabled, all workstations are allowed to access the ProxySG administration points. You can also add allowed workstations later to the access control list (ACL). (For more information on limiting workstation access, see ["Moderate Security: Restricting Management Console Access Through the Console Access Control List \(ACL\)"](#) on page 63.)

Securing the Serial Port

If you choose to secure the serial port, you must provide a Setup Console password that is required to access the Setup Console in the future.

Once the secure serial port is enabled:

- ❑ The Setup Console password is required to access the Setup Console.
- ❑ An authentication challenge (username and password) is issued to access the CLI through the serial port.

To recover from a lost Setup Console password, you can:

- ❑ Use the Front Panel display to either disable the secure serial port or enter a new Setup Console password.
- ❑ Use the CLI `restore-defaults factory-defaults` command to delete all system settings. For information on using the `restore-defaults factory-defaults` command, see ["Factory-Defaults"](#) on page 1403.
- ❑ Use the reset button (if the appliance has a reset button) to delete all system settings. Otherwise, reset the ProxySG to its factory settings by holding down the left arrow key on the front-panel for 5 seconds. The appliance will be reinitialized.

To reconfigure the appliance or secure the serial port, refer to the [hardware guides](#) for your appliance.

About Password Security

In the ProxySG, the console administrator password, the Setup Console password, and Enable (privileged-mode) password are hashed and stored. It is not possible to reverse the hash to recover the plain text passwords.

In addition, the `show config` and `show security` CLI commands display these passwords in their hashed form. The length of the hashed password depends on the hash algorithm used so it is not a fixed length across the board.

Passwords that the ProxySG uses to authenticate itself to outside services are encrypted using triple-DES on the appliance, and using RSA public key encryption for output with the `show config` CLI command. You can use a third-party encryption application to create encrypted passwords and copy them into the ProxySG using an encrypted-password command (which is available in several modes and described in those modes). If you use a third-party encryption application, verify it supports RSA encryption, OAEP padding, and Base64 encoded with no new lines.

These passwords, set up during configuration of the external service, include:

- ❑ Access log FTP client passwords (primary, alternate)—For configuration information, see ["Editing the FTP Client"](#) on page 636.
- ❑ Archive configuration FTP password—For configuration information, see [Chapter 5: "Backing Up the Configuration"](#) on page 69.
- ❑ RADIUS primary and alternate secret—For configuration information, see [Chapter 58: "RADIUS Realm Authentication and Authorization"](#) on page 1089.
- ❑ LDAP search password—For configuration information, see ["Defining LDAP Search & Group Properties"](#) on page 1052.
- ❑ Content filter download passwords—For configuration information, see ["Downloading a Content Filter Database"](#) on page 395.

Limiting User Access to the ProxySG—Overview

When deciding how to give other users read-only or read-write access to the ProxySG, sharing the basic console account settings is only one option. The following summarizes all available options:

Note: If Telnet Console access is configured, Telnet can be used to manage the ProxySG with behavior similar to SSH with password authentication.

SSL configuration is not allowed through Telnet, but is permissible through SSH.

Behavior in the following sections that applies to SSH with password authentication also applies to Telnet. Use of Telnet is not recommended because it is not a secure protocol.

- ❑ Console account—minimum security

The console account username and password are evaluated when the ProxySG is accessed from the Management Console through a browser and from the CLI through SSH with password authentication. The Enable (privileged-mode) password is evaluated when the console account is used through SSH with password authentication and when the CLI is accessed through the serial console and through SSH with RSA authentication. The simplest way to give access to others is sharing this basic console account information, but it is the least secure and is not recommended.

To give read-only access to the CLI, do not give out the Enable (privileged-mode) password.

❑ Console access control list—moderate security

Using the access control list (ACL) allows you to further restrict use of the console account and SSH with RSA authentication to workstations identified by their IP address and subnet mask. When the ACL is enforced, the console account can only be used by workstations defined in the console ACL. Also, SSH with RSA authentication connections are only valid from workstations specified in the console ACL (provided it is enabled).

After setting the console account username, password, and Enable (privileged-mode) password, use the CLI or the Management Console to create a console ACL. See "[Moderate Security: Restricting Management Console Access Through the Console Access Control List \(ACL\)](#)" on page 63.

❑ Per-user RSA public key authentication—moderate security

Each administrator's public keys are stored on the appliance. When connecting through SSH, the administrator logs in with no password exchange. Authentication occurs by verifying knowledge of the corresponding private key. This is secure because the passwords never go over the network.

This is a less flexible option than CPL because you cannot control level of access with policy, but it is a better choice than sharing the console credentials.

❑ Blue Coat Content Policy Language (CPL)—maximum security

CPL allows you to control administrative access to the ProxySG through policy. If the credentials supplied are not the console account username and password, policy is evaluated when the ProxySG is accessed through SSH with password authentication or the Management Console. Policy is never evaluated on direct serial console connections or SSH connections using RSA authentication.

- Using the CLI or the Management Console GUI, create an authentication realm to be used for authorizing administrative access. For administrative access, the realm must support BASIC credentials—for example, LDAP, RADIUS, Local, or IWA with BASIC credentials enabled.
- Using the Visual Policy Manager, or by adding CPL rules to the Local or Central policy file, specify policy rules that: (1) require administrators to log in using credentials from the previously-created administrative realm, and (2) specify the conditions under which administrators are either denied all access, given read-only access, or given read-write access. Authorization can be based on IP address, group membership, time of day, and many other conditions. For more information, refer to the *Visual Policy Manager Reference*.
- To prevent anyone from using the console credentials to manage the ProxySG, set the console ACL to deny all access (unless you plan to use SSH with RSA authentication). For more information, see "[Moderate Security: Restricting Management Console Access Through the Console Access Control List \(ACL\)](#)" on page 63. You can also restrict access to a single IP address that can be used as the emergency recovery workstation.

The following chart details the various ways administrators can access the ProxySG console and the authentication and authorization methods that apply to each.

Table 4–1 ProxySG Console Access Methods/Available Security Measures

Security Measures Available	Serial Console	SSH with Password Authentication	SSH with RSA Authentication	Management Console
Username and password evaluated (console-level credentials)		X		X
Console Access List evaluated		X (if console credentials are offered)	X	X (if console credentials are offered)
CPL <Admin> Layer evaluated		X (see Note 1 below)		X (see Note 2 below)
Enable password required to enter privileged mode (see Note 2 below)	X	X	X	
CLI line-vty timeout command applies.	X	X	X	
Management Console Login/Logout				X

Notes


- ❑ When using SSH (with a password) and credentials other than the console account, the enable password is actually the same as the login password. The privileged mode password set during configuration is used only in the serial console, SSH with RSA authentication, or when logging in with the console account.
- ❑ In this case, user credentials are evaluated against the policy before executing each CLI command. If you log in using the console account, user credentials are not evaluated against the policy.

Moderate Security: Restricting Management Console Access Through the Console Access Control List (ACL)

The ProxySG allows you to limit access to the Management Console and CLI through the console ACL. An ACL, once set up, is enforced only when console credentials are used to access either the CLI or the Management Console, or when an SSH with RSA authentication connection is attempted. The following procedure specifies an ACL that lists the IP addresses permitted access.

To create an ACL:

1. Select **Configuration > Authentication > Console Access > Console Access**.



2. (Optional) Add a new address to the ACL:
 - a. Click **New**. The Add List Item dialog displays.
 - b. In the **IP/Subnet** fields, enter a static IP address. In the **Mask** fields, enter the subnet mask. To restrict access to an individual workstation, enter 255.255.255.255.
 - c. Click **OK** to add the workstation to the ACL and return to the **Console Access** tab.
3. Repeat step 2 to add other IP addresses.
4. To impose the ACL defined in the list box, select **Enforce ACL for built-in administration**. To allow access to the CLI or Management Console using console account credentials from any workstation, clear the option. The ACL is ignored.

Important: Before you enforce the ACL, verify the IP address for the workstation you are using is included in the list. If you forget, or you find that you mis-typed the IP address, you must correct the problem using the serial console.

5. Click **Apply**.

Maximum Security: Administrative Authentication and Authorization Policy

The ProxySG permits you to define a rule-based administrative access policy. This policy is enforced when accessing:

- ❑ the Management Console through HTTP or HTTPS

- ❑ the CLI through SSH when using password authentication
- ❑ the CLI through telnet
- ❑ the CLI through the serial port if the secure serial port is enabled

These policy rules can be specified either by using the VPM or by editing the Local policy file. Using policy rules, you can deny access, allow access without providing credentials, or require administrators to identify themselves by entering a username and password. If access is allowed, you can specify whether read-only or read-write access is given. You can make this policy contingent on IP address, time of day, group membership (if credentials were required), and many other conditions.

Serial-console access is not controlled by policy rules. For maximum security to the serial console, physical access must be limited.

SSH with RSA authentication also is not controlled by policy rules. You can configure several settings that control access: the enable password, the console ACL, and per-user keys configured through the **Configuration > Services > SSH > SSH Client** page. (If you use the CLI, SSH commands are under **Configuration > Services > SSH-Console**.)

Defining Administrator Authentication and Authorization Policies

Administrative authentication uses policy, (either Visual Policy or CPL in the local policy file) to authenticate administrative users to the appliance. This is done with two layers in policy: one to define the realm that is used to authenticate users (**Admin Authentication layer**) and the other to define security rights for authenticated users or groups (**Admin Access layer**).

Note: If you choose a realm that relies on an external server and that server is unavailable, the appliance will not be able to authenticate against that realm.

For best security, the following authentication realms are recommended by Symantec for administrative authentication to the appliance.

- ❑ IWA-BCAAA (with TLS -- not SSL) with basic credentials
- ❑ Local
- ❑ .509 certificate based (including certificate realms; refer to the *Common Access Card Solutions Guide* for information)
- ❑ LDAP with TLS (not SSL)
- ❑ IWA-Direct with basic credentials
- ❑ RADIUS

The following realms can be configured for administrative authentication, but pass administrative credentials in clear text. These realms should not be used for administrative authentication:

- ❑ Windows SSO

- ❑ Novell SSO
- ❑ IWA-BCAAA without SSL or TLS
- ❑ LDAP without SSL or TLS

The following realms do not support administrative authentication:

- ❑ IWA-BCAAA/IWA-Direct realms that do not accept basic credentials
- ❑ SiteMinder
- ❑ COREid
- ❑ SAML (Policy Substitution)
- ❑ XML

Note: Other authentication realms can be used, but will result in administrative credentials being sent in clear text.

Configure Administrative Authentication with a Local Realm

The process to provide read-only access for administrators includes the following steps:

- ❑ Create a local authentication realm.
- ❑ Create a list that includes usernames and passwords for members whom you wish to provide read-only access in the Management Console.
- ❑ Connect the list to the local realm.
- ❑ Create policy to enforce read-only access to members included in the list.


Use the steps below to complete the tasks detailed above.

1. Create a local realm:
 - a. Select the **Configuration > Authentication > Local > Local Realms** tab.
 - b. Click **New** to add a new realm. In this example the realm is named **MC_Access**.
2. Using the Command Line Interface (CLI), create a list of users who need read-only access. The list must include a username and password for each user.
 - a. Enter configuration mode in the CLI; this example creates a list called `Read_Access`.


```
SGOS#(config)security local-user-list create Read_Access
```
 - b. Edit the list to add user(s) and to create usernames and passwords. This example adds a user named `Bob Kent`.

```
SGOS#(config)security local-user-list edit Read_Access
SGOS#(config)user create Bob_Kent
SGOS#(config)user edit Bob_Kent
SGOS#(config)password 12345
```



3. Connect the user list (created in Step 2) to the local realm (created in Step 1).
 - a. In the **Configuration > Authentication > Local > Local Main** tab, select **MC_Access** from the **Realm name** drop-down menu.
 - b. Select **Read_Access** from the **Local user list** drop-down menu.
4. Use the for creating policy to enforce read-only access to the users in your list:
 - a. Launch the VPM.
 - b. Create an **Admin Authentication Layer** (or add a new rule in an existing layer). This layer determines the authentication realm that will be used to authenticate users who access the Management Console of the ProxySG.



- c. In the **Action** column, right click and select **Set**. In the **Set Action** dialog that displays, click **New** and select **Authenticate**. The Add Authentication Object displays.



- d. In the **Add Authenticate Object** dialog that displays, select the local realm you created in Step 1.
 - e. Create an **Admin Access Layer**.
 - f. In the **Source** column, right click and select **Set**. In the **Set Source Object** dialog that displays, click **New** and select **User**. The Add User Object dialog displays.



- g. Enter the name of the user for whom you want to provide read-only access.
- h. Click **OK** in both dialogs.



- i. In the **Action** column, right click and select **Allow Read-only Access**.
5. Click **Install Policy**.

The user can now log in the Management Console as a user with read-only access. Repeat step 4 and use Allow Read/Write access to define user access with read/write privileges

Chapter 5: Backing Up the Configuration

This chapter describes how to back up your configuration and save it on a remote system so that you can restore it in the unlikely event of system failure or replacement. ProxySG configuration backups are called *archives*.

Important: You should archive the system configuration before performing any software or hardware upgrade or downgrade.

System archives can be used to

- ❑ Restore the appliance to its previous state in case of error.
- ❑ Restore the appliance to its previous state because you are performing maintenance that requires a complete restoration of the system configuration. For example, upgrading all the disk drives in a system.
- ❑ Save the system configuration so that it can be restored on a replacement appliance. This type of configuration archive is called a *transferable* archive.
- ❑ Propagate configuration settings to newly-manufactured ProxySG appliances. This process is called *configuration sharing*.

Topics in this Chapter

The following topics are covered in this chapter:

- ❑ [Section A: "About Configuration Archives" on page 70](#)
- ❑ [Section B: "Archiving Quick Reference" on page 72](#)
- ❑ [Section C: "Creating and Saving a Standard Configuration Archive" on page 75](#)
- ❑ [Section D: "Creating and Saving a Secure \(Signed\) Archive" on page 77](#)
- ❑ [Section E: "Preparing Archives for Restoration on New Devices" on page 80](#)
- ❑ [Section F: "Uploading Archives to a Remote Server" on page 90](#)
- ❑ [Section G: "Restoring a Configuration Archive" on page 94](#)
- ❑ [Section H: "Sharing Configurations" on page 96](#)
- ❑ [Section I: "Troubleshooting" on page 98](#)

Section A: About Configuration Archives

This section describes the archive types and explains archive security and portability.

This section includes the following topics:

- ❑ ["About the Archive Types and Saved Information"](#) on page 70
- ❑ ["About Archive Security"](#) on page 70
- ❑ ["About Archive Portability"](#) on page 71
- ❑ ["What is not Saved"](#) on page 71

About the Archive Types and Saved Information

Three different archive types are available. Each archive type contains a different set of configuration data:

- ❑ **Configuration - post setup:** This archive contains the configuration on the current system—minus any configurations created through the setup console, such as the IP address. It also includes the installable lists but does not include SSL private key data. Use this archive type to share an appliance's configuration with another. See ["Sharing Configurations"](#) on page 96 for more information.
- ❑ **Configuration - brief:** This archive contains the configuration on the current system and includes the setup console configuration data, but does not include the installable lists or SSL private key and static route information.

Note: An installable list is a list of configuration parameters that can be created through a text editor or through the CLI inline commands and downloaded to the ProxySG from an HTTP server or locally from your PC. Configurations that can be created and installed this way include ProxyClient, archiving, forwarding hosts, SOCKS gateways, policy files, and exceptions.

- ❑ **Configuration - expanded:** This is the most complete archive of the system configuration, but it contains system-specific settings that might not be appropriate if pushed to a new system. It also does not include SSL private key data. If you are trying to create the most comprehensive archive, Blue Coat recommends that you use the configuration-expanded archive.

Options in the Management Console enable you to create standard, secure, and transferable versions of the three archive types.

About Archive Security

The ProxySG provides two methods for creating archives, *signed* and *unsigned*. A signed archive is one that is cryptographically signed with a key known only to the signing entity—the digital signature guarantees the integrity of the content and the identity of the originating device.

To create signed archives, your appliance must have an SSL certificate guaranteed by a CA. You can then use a trusted CA Certificate List (CCL) to verify the authenticity of the archive.

Use signed archives only when security is high priority. Otherwise, use unsigned archives. For information about creating secure archives, see ["Creating and Saving a Secure \(Signed\) Archive"](#) on page 77.

About Archive Portability

To retain the option to transfer the configuration from the source appliance to another appliance, the configuration cannot be restored unless you save the SSL keyrings, and the `configuration-passwords-key` in particular.

The `configuration-passwords-key` keyring must be saved. This keyring is used to encrypt and decrypt the passwords (login, enable, FTP, etc.) and the passwords cannot be restored without it. This is because the purpose of public/private key authentication is to disallow decryption by a device other than the device with the private key. To restore any encrypted data from an archive, you must have the corresponding SSL keyring.

See ["Creating a Transferable Archive"](#) on page 81 for more information about creating transferable archives.

What is not Saved

Archiving saves the ProxySG appliance configuration only. Archives do not save the following:

- ❑ Cache objects
- ❑ Access logs
- ❑ Event logs
- ❑ License data (you might need to reapply the licenses)
- ❑ Software image versions
- ❑ SSL key data
- ❑ Content-filtering databases
- ❑ Exception pages

Section B: Archiving Quick Reference

This section provides a table of quick reference tasks and describes the high-level archive creation and restoration tasks.

This section includes the following topics:

- ❑ "Archiving Quick Reference Table" on page 72
- ❑ "Overview of Archive Creation and Restoration" on page 73

Archiving Quick Reference Table

The following table lists common archive management tasks and where to get more information.

Table 5–1 Archiving Task Table

If You Want to...	Go To...
Understand the archive and restoration process	"Overview of Archive Creation and Restoration" on page 73
Find out what is not archived	"What is not Saved" on page 71
Learn about the archive types	"About the Archive Types and Saved Information" on page 70
Learn about secure archives	"About Archive Security" on page 70
Learn about transferable archives A transferable archive is a configuration archive that can be imported to a new device.	"About Archive Portability" on page 71
Create a standard archive	"Creating and Saving a Standard Configuration Archive" on page 75
Create a secure archive	"Creating and Saving a Secure (Signed) Archive" on page 77
Create a transferable archive	"Creating a Transferable Archive" on page 81
Upload an archive to a remote server	"Uploading Archives to a Remote Server" on page 90
Schedule archive creation	You cannot schedule archive creation from the ProxySG appliance; you must use Director for that. Refer to "Creating, Scheduling, and Managing Jobs" in the <i>Director Configuration and Management Guide</i> .
Understand file name identifiers	"Adding Identifier Information to Archive Filenames" on page 92

Table 5–1 Archiving Task Table (Continued)

If You Want to...	Go To...
Restore an archive	"To install the archived configuration:" on page 94
Share Configurations	"Sharing Configurations" on page 96
Troubleshoot archive configuration	"Troubleshooting" on page 98

Overview of Archive Creation and Restoration

The following list describes all of the possible steps required to create and restore an unsigned, signed, or transferable configuration archive. You do not have to perform all of these steps to complete a standard, unsigned archive. Non-standard archiving steps are indicated by the word "Optional."

1. Optional (for transferable archives only)—Record the `configuration-passwords-key` data on the source ProxySG, as described in "Option 1: Recording SSL Keyring and Key Pair Information" on page 82. If you need to restore the archive onto a different appliance, you must have this data.

Do not lose the password used to encrypt the private key. If you do, you will not be able to recover your private keys.

2. Optional (for transferable archives only)—Record any other SSL keyring data you want to save.
3. Determine the type of archive to create—secure or standard. See "About Archive Security" on page 70.

If you are creating an standard archive, go to Step 5. Otherwise, go to Step 4.

4. Optional (for secure archives only)—Verify that the source ProxySG has an appliance certificate, as described in "Using the Appliance Certificate to Sign the Archive" on page 77. If it does not have an appliance certificate:
 - a. Create a keyring on the appliance.
A keyring contains a public/private key pair. It can also contain a certificate signing request or a signed certificate.
 - b. Create a Certificate Signing Request (CSR) and send it to a Certificate Signing Authority (CA).
 - c. Have the CA sign the CSR.

To get more information about appliance certificates, see Chapter 61: "Managing X.509 Certificates".

5. Archive the configuration:
 - Standard, unsigned archive—"Creating and Saving a Standard Configuration Archive" on page 75.
 - Secure archive—"Creating and Saving a Secure (Signed) Archive" on page 77

- Transferable archive—"Creating a Transferable Archive" on page 81.
6. Store the archive in a secure location.
 7. If you are restoring the archive to another device, import the `configuration-passwords-key` onto the target device, as described in "Restoring an Archived Key Ring and Certificate" on page 88.
 8. Restore the archive, as described in "Restoring a Configuration Archive" on page 94.

Figure 5–1 on page 74 describes the archive creation process.





Figure 5–1 Flow Chart of Archive Creation Process

Section C: Creating and Saving a Standard Configuration Archive

Use the Management Console to create a standard archive of the system configuration. This is the simplest method of archive creation. This type of archive cannot be transferred to another appliance unless you save the SSL keyrings as described in [Section E: "Preparing Archives for Restoration on New Devices"](#) on page 80.

To create a standard configuration archive:

1. Access the Management Console of the ProxySG you want to back up:
`https://ProxySG_IP:8082`
2. Select **Configuration > General > Archive**. The **Archive Configuration** tab displays.



3. Select a configuration type:
 - a. In the **View Current Configuration** section, select **Configuration - expanded** from the View File drop-down list.
 - b. View the configuration you selected by clicking **View**.
 A browser window opens and displays the configuration.

Note: You can also view the file by selecting **Text Editor** in the **Install Configuration** panel and clicking **Install**.

4. Save the configuration.
You can save the file two ways:
 - Use the browser **Save As** function to save the configuration as a text file on your local system. This is advised if you want to re-use the file.
 - Copy the contents of the configuration. (You will paste the file into the Text Editor on the newly-manufactured system.)

To restore a standard archive:

1. Select **Configuration > General > Archive**.
2. Select **Local File** and click **Install**.

3. Browse to the location of the archive and click **Open**. The configuration is installed, and the results screen displays.

Section D: Creating and Saving a Secure (Signed) Archive

This section describes how to use the Management Console to save a secure (signed) archive of the system configuration. A signed archive is an archive signed with a digital signature that can only be read by the device that created it, thus guaranteeing the integrity and authenticity of the archive. To create signed archives, your appliance must have an SSL certificate guaranteed by a CA.

Signed archives have a `.bcsc` extension and contain the following files:

- ❑ `show configuration output`
- ❑ PKCS#7 detached signature

This section includes the following topics:

- ❑ ["Using the Appliance Certificate to Sign the Archive"](#) on page 77
- ❑ ["Creating Signed Configuration Archives"](#) on page 78
- ❑ ["Modifying Signed Archives"](#) on page 79

Before Reading Further

If you are not familiar with SSL authentication, read the following before proceeding:

- ❑ ["About Archive Security"](#) on page 70
- ❑ The device authentication information in [Chapter 71: "Authenticating a ProxySG"](#) on page 1289.
- ❑ The X.509, CCL, and SSL information in [Chapter 61: "Managing X.509 Certificates"](#).

Using the Appliance Certificate to Sign the Archive

If your appliance has a built-in appliance certificate, you can use it, and the corresponding `appliance-ccl` CCL, to sign the archive.

Note: ProxySG appliances manufactured before July 2006 do not support appliance certificates. These devices must use an SSL certificate guaranteed by a CA to sign archives, as described below.

To determine if your device has an appliance certificate:

1. Use an SSH client to establish a CLI session with the ProxySG.
2. Enter enable mode:
`SGOS # enable`
3. Enter the following command:
`SGOS # show ssl certificate appliance-key`

The appliance certificate displays if the appliance has one. Otherwise, the following error is displayed:

```
Certificate "appliance-key" not found
```

4. If the appliance does not have an appliance certificate, create one as follows:

- a. Create a keyring on the appliance.

A keyring contains a public/private key pair. It can also contain a certificate signing request or a signed certificate.

- b. Create a Certificate Signing Request (CSR) and send it to a Certificate Signing Authority (CA).
- c. Have the CA sign the CSR (this process results in a digital certificate).
- d. Import the keyring and certificate as described in ["Restoring an Archived Key Ring and Certificate"](#) on page 88.

For more information about appliance certificates, see [Chapter 61: "Managing X.509 Certificates"](#).

Creating Signed Configuration Archives


This section describes how to save a signed configuration archive to the computer you are using to access the Management Console.

To create and save a signed configuration archive to your computer:

1. Access the Management Console of the ProxySG you want to back up:

```
https://ProxySG_IP:8082
```

2. Select the **Configuration > General > Archive > Archive Storage** tab.



3. From the **Sign archives with keyring** drop-down list, select a signing keyring to use or accept the default (**appliance-key**).
4. Click **Apply**.

Note: If you do not click **Apply**, a pop-up displays when you click **Save** that indicates that all unsaved changes will be saved before storing the archive configuration. The unsaved changes are the **Sign archives with keyring** option changes you made in Step 3.

5. From the **Save archive** drop-down list, select the archive type (Blue Coat recommends **Configuration - expanded**).
6. Click **Save**.

A new browser window displays, prompting you to open or save the configuration to the local disk of the device you are using to access the ProxySG.

To restore a signed archive:

1. Connect to the appliance Management Console of the target appliance, that is the ProxySG that you are installing the configuration onto.

`https://ProxySG_IP:8082`

2. Go to the Management Console Home page and view the **Software version:** information to verify that the appliance is running the same software version that was used to create the archive. For example:

Software version: SGOS 5.3.0.2 Proxy Edition

You can also verify the version from the appliance CLI:

```
SGOS # enable
SGOS # show version
```

3. Select **Configuration > General > Archive**.
4. In the **Install Configuration** panel, check the setting of the **Enforce installation of signed archives** option. If this option is selected, only signed archives can be restored.
5. Select a CCL to use to verify the archive from the **Verify signed archive with CCL** drop-down list. If you used the **appliance-key** keyring, select **appliance-ccl**.
6. Select **Local File** and click **Install**.

Modifying Signed Archives

If you modify a signed archive, you must subsequently restore it as an unsigned archive.

If you created a signed archive and want to verify its authenticity before modifying it, use OpenSSL or another tool to verify the signature before making modifications. (The use of OpenSSL is beyond the scope of this document.) Because a signed archive contains the output of the `show configuration` command, you can extract the `show configuration` command output, modify it as required, and treat the archive as unsigned thereafter.

Section E: Preparing Archives for Restoration on New Devices

While a configuration archive will back up the appliance configuration, that configuration cannot be transferred to another device unless you save the SSL keyrings on the appliance—especially the `configuration-passwords-key` keyring. The process of creating the archive and saving the associated SSL keyrings is called creating a *transferable archive*.

Note: You must also save the SSL keyrings if you plan to restore an encrypted archive after a reinitialization. When you reinitialize the appliance, new keys get created, and you will therefore not be able to restore the configuration unless you first restore the `configuration-passwords-key`.

This section includes the following topics:

- ❑ ["About the configuration-passwords-key" on page 80](#)
- ❑ ["Creating a Transferable Archive" on page 81](#)
- ❑ ["Option 1: Recording SSL Keyring and Key Pair Information" on page 82](#)
- ❑ ["Option 2: Changing Encrypted Passwords to Clear Text" on page 87](#)
- ❑ ["Restoring an Archived Key Ring and Certificate" on page 88](#)

About the configuration-passwords-key

The `configuration-passwords-key` is an SSL keyring. SSL is a method of securing communication between devices. SSL uses a public key to encrypt data and private key to decrypt data. These keys (stored in “keyrings”) are unique to the device. This ensures that data encrypted with a device’s public key can only be decrypted by the corresponding private key.

On ProxySG appliances, the `configuration-passwords-key` SSL keyring is used to encrypt and decrypt the following passwords on the appliance:

- ❑ Administrator console passwords (not needed for shared configurations)
- ❑ Privileged-mode (enable) passwords (not needed for shared configurations)
- ❑ The front-panel PIN (recommended for limiting physical access to the system)
- ❑ Failover group secret
- ❑ Access log FTP client passwords (primary, alternate)
- ❑ Archive configuration FTP password
- ❑ RADIUS primary and alternate secret
- ❑ LDAP search password
- ❑ SNMP read, write, and trap community strings
- ❑ RADIUS and TACACS+ secrets for splash pages

Because every appliance has a different `configuration-passwords-key`, you will receive a decryption error if you try to restore an archive to another device.

To ensure that the archive can be transferred to another appliance, you must do one of the following:

- ❑ Restore the original `configuration-passwords-key` keyring

While it is possible to reset each of the passwords using the Management Console, it is easier to save the original keyring so that you can import it to the new appliance (before restoring the configuration). Restoring the keyring allows all previously configured passwords to remain valid after archive restoration.

- ❑ Change the encrypted passwords to clear text so that they can be regenerated.

Note: To save an SSL keyring, you must be able to view it. If the key is marked `no-show`, you cannot save it.

Creating a Transferable Archive

This section describes the steps required to create a transferable archive.

To create a transferable archive:

1. Record the `configuration-passwords-key` data on the source ProxySG, as described in "[Option 1: Recording SSL Keyring and Key Pair Information](#)" on page 82. If you need to restore the archive onto a different appliance, you must have this data.

Do not lose the password used to encrypt the private key. If you do, you will not be able to recover your private keys.

2. Record any other SSL keyring data you want to save.
3. Store the keyring data and archive in a secure location.
4. Create the archive as described in "[Creating and Saving a Standard Configuration Archive](#)" on page 75.

To restore a transferable archive:

1. Connect to the appliance Management Console of the target appliance, that is the ProxySG that you are installing the configuration onto.

`https://ProxySG_IP:8082`

2. Go to the Management Console Home page and view the **Software version:** information to verify that the appliance is running the same software version that was used to create the archive. For example:

Software version: SGOS 5.3.0.2 Proxy Edition

You can also verify the version from the appliance CLI:

```
SGOS # enable
SGOS # show version
```

3. Restore the `configuration-passwords-key` data and any other SSL key data.
Import the `configuration-passwords-key` keyring as described in ["Restoring an Archived Key Ring and Certificate"](#) on page 88.
4. Select **Configuration > General > Archive**.
5. Select **Local File** and click **Install**.
6. Browse to the location of the archive and click **Open**. The configuration is installed, and the results screen displays.

Option 1: Recording SSL Keyring and Key Pair Information

For security reasons, Blue Coat recommends that you *do not* change encrypted passwords to clear text. Instead, preserve the `configuration-passwords-key` keyring on the source device (the appliance that you created the archive from) and import that keyring to the target device before you restore the archive.

You can also use the following procedure to save any other keyrings required to reload SSL-related configuration that references those keyrings.

To record the `configuration-passwords-key` keyring on the source ProxySG:

1. Copy the following template to a text file and use it to record the certificate information so that you can import and restore it later. This template allows you to import a certificate chain containing multiple certificates, from the CLI.
Alternatively, you can simply copy the SSL data into a blank text file.

Note: The following example is shown in smaller text to preserve the structure of the commands.

```
!  
ssl ; switches from config mode to config ssl  
!  
inline keyring show configuration-passwords-key "end-inline"  
!  
end-inline  
inline keyring show default "end-inline"  
!  
end-inline  
!  
inline certificate default "end-inline"  
!  
end-inline  
!  
! repeat this process for each keyring. Be sure to import the private  
key first, then the keyrings certificate  
!
```



```
exit ; returns to config mode
!
```

Do not specify your passwords; the system will prompt you for them when you restore the keys (SGOS 5.3 and later). You can modify the template to include other keyrings and certificates.

2. From the CLI, access the `config` prompt (using the serial console or SSH):

```
sgos # config terminal
```

3. Enter the following commands:

```
sgos #(config) ssl
```

```
sgos #(config ssl) view keyring
```

A listing of existing keyrings (and certificates) is displayed.

For example (your keyrings might be different):

```
sgos #(config ssl) view keyring
```

```
Keyring ID:          appliance-key
Private key showability: no-show
Signing request:     present
Certificate:          absent
```

```
Keyring ID:          configuration-passwords-key
Private key showability: show
Signing request:     absent
Certificate:          absent
```

```
Keyring ID:          default
Private key showability: show
Signing request:     absent
Certificate:          present
Certificate issuer:   Blue Coat SG200 Series
Certificate valid from: Dec 04 20:11:04 2007 GMT
Certificate valid to:   Dec 03 20:11:04 2009 GMT
Certificate thumbprint:
9D:B2:36:E5:3D:B7:88:21:CB:0A:08:39:2C:A1:4B:CB
```

```
Keyring ID:          passive-attack-protection-only-key
Private key showability: show
Signing request:     absent
Certificate:          present
Certificate issuer:   Blue Coat SG200 Series
Certificate valid from: Dec 04 20:11:07 2007 GMT
Certificate valid to:   Dec 03 20:11:07 2009 GMT
Certificate thumbprint:
0B:AD:07:A7:CF:D9:58:03:89:5B:67:35:43:B9:F2:C9
```

4. Enter the following command:

```
sgos #(config ssl) view keypair des3 configuration-passwords-key
```

Note: The `aes128` and `aes256` encryption options are also supported.

- When prompted, enter an encryption key password:

```
Encryption key: *****
Confirm encryption key: *****
```

This password is used to encrypt the private-key before displaying it. After confirming the password, the ProxySG displays the encrypted private-key associated with that keyring.

Important: Do not lose the password used to encrypt the private key. If you do, you will not be able to recover your private keys.

For example:

```
sgos #(config ssl)view keypair des3 configuration-passwords-key
Encryption password: *****
Confirm encryption password: *****
-----BEGIN RSA PRIVATE KEY-----
Proc-Type: 4,ENCRYPTED
DEK-Info: DES-EDE3-CBC,D542F10E3FFF899F

aFqxQNOD+321IXdQjCGmT+adeQqMiQDAyCOvWd+aJ+OmDjITpd7bwijcxWA89RB8
y65NSia0UmTClY9MM4j6T/fXhBspEu7Wyc/nM+005pJldxTmZgPig6TiIiOlXtMI
ymCLolxjAr+vFSx7ji6jUT13JxZHfksNd9DS06DHLr6hJNERDi9dGog561zlwBo8
zvs0x4PqB+mQ05qewmReMs9tnuLkGgBXguH+2Nw9hI0WKEa9KPFWRznD/+zEZbEo
nM+VOwn3nWuqcFRLFoSUP2QBZ581pU3XAUydabBn0uBOMR4a3C+F/W/v0p71jJ9o
JL6Ao/S46A4UgPkuswGMXo1kG3K2J/Ev4nMBua6HSZgM87DxvMSiCZ1XxlKlBqv
F9P+1l03mdR3g2LzK1DLTvlcA9pEPbW65gmnpgj/WLqhEyNpm+DkplxMtMESxNqM
4attb8fXAEcRI+liUWpjxnycqlm+dcFqq6/bLixYSQ4HGXFx5qTot+FtIvB5h3g
KwQusgaLVTiesn9K7BQK4wjXJK1DclIrog+ET1fkxtj2oA5/7HN10Ar0ogBxsZLj
0LS5fwVfHNkuyNLUXZSAiLlLoIqFIVtRiRfiWe3e/eJvazIaErEk40NvIaaXP1j9p
ENzK2dw9WS7xtcU5kAcdoiX1lFONauKDVUkHwhvqz3KnMt1p81fkDUpiD1xaVfMg
s2FapgjAsYciEJxDUfPLzYVlvpOpx6DW3t0D0AlEKkPVNm9RzlnXjk2CPTdPERC
pKN+EIKS2kqpRE6hHu37zzN06ipPNu2cCSHI/ozc0X4=
-----END RSA PRIVATE KEY-----
```

- Copy the configuration-passwords-key and paste it into the template (copied in step 1) beneath the line inline keyring show configuration-passwords-key "end-inline".
- If a certificate is associated with a keyring, enter the following command:

```
sgos #(config ssl) view certificate keyring-name
```

For example:

```
sgos #(config ssl)view certificate appliance-key
-----BEGIN CERTIFICATE-----
MIICUzCCAbygAwIBAgIEFm6QWzANBgkqhkiG9w0BAQUFADBUMQswCQYDVQQGDAIg
IDETMBEGA1UECAwKU29tZS1tdGF0ZTEFMBOGA1UECgwWQmx1ZSBDb2F0IFNBMjAw
IFNlcmllczETMBEGA1UECwwKNDYwNTA2MDAwMTEUMBIGA1UEAwwLMTAuOS41OS4y
```

```
MTAwHhcNMDcxMjA0MjAxMTA3WhcNMDkxMjAzMjAxMTA3WjBuMQswCQYDVQQGDAIg
IDETMBEGA1UECAwKU29tZS1TdGF0ZTEfMB0GA1UECgwWQmx1ZSBDb2F0IFNBMjAw
IFNlcmllczETMBEGA1UECwwKNDYwNTA2MDAwMTEUMBIGAlUEAwLMTAuOS41OS4y
MTAwgZ8wDQYJKoZIhvcNAQEBBQADgY0AMIGJAoGBAJ/F/Sn3CzYvbFPWDD03g9Y/
O3jwCrcXLU8cki6SZUVl9blgzBTgBY3KyDl2baqZNl2QGwkspEtDI45G3/K2GRIF
REs3mKGxY7fbwgRpoL+nRT8w9qWHO393pGr1JKFldXbYOzn3p31EXUuGRfXkIqeA
919uvOD5gOX0BEzrvDRnAgMBAAEwDQYJKoZIhvcNAQEFBQADgYEASgIR9r2MuRbc
ltHq/Lb5rIXn13wFZEND/viO54YOiW1ZixlpCBbDIkef3DdJZLxVy3x7GbW32OfE
3a7kfIMvVKWmNO+syAn4B2yasy0nxbSyOciJq1C42yPJ+Bj1MuYDmgIvMP6ne5UA
gYYhe/koamOZNcIuaXrAS2v2tYevrBc=
-----END CERTIFICATE-----
```

8. Copy the certificate and paste it into the template (copied in step 1) beneath the inline certificate `cert_name "end-inline" line`.
9. Optional—For *each* named keyring that you want to restore, repeat steps 4 to 8.

Note: The appliance-key keyring's private key is not viewable, and cannot be transferred to another ProxySG. The default and passive-attack-protection-only-key keys typically do not need to be restored either.

10. Save the template with the configuration-passwords-key and other SSL key data on a secure server.
11. Save the password information (that you used to encrypt the keys) in a secure place, for example, a restricted access cabinet or safe.

After saving this data, create a configuration archive as described in ["Creating a Transferable Archive"](#) on page 81. When you are ready to restore the archive, you must first restore the SSL data on the target appliance as described in ["Restoring an Archived Key Ring and Certificate"](#) on page 88.

Example: Completed SSL Data Template

The following example shows how the template might look after completing the procedure in ["To record the configuration-passwords-key keyring on the source ProxySG:"](#) on page 82.

The template allows you to import a certificate chain containing multiple certificates, from the CLI. When you restore the data to the appliance, you will be prompted for the encryption password that you used to encrypt the keys.

Note: The commands in the following example are bounded by the document text area and wrap to the next line. They are not shown here as they would appear in the CLI. See Step 1 in ["Option 1: Recording SSL Keyring and Key Pair Information"](#) on page 82 to view an example of how the commands should appear.

!

```
ssl ; switches from config mode to config ssl
!
inline keyring show configuration-passwords-key "end-inline"
-----BEGIN RSA PRIVATE KEY-----
Proc-Type: 4,ENCRYPTED
DEK-Info: DES-EDE3-CBC,2F6148C8A9902D7F

11JjGKxpkcWBXj424FhyQJPKRdgHUIx12C6HKigth6hUgPqsSJj958FbzEx6ntsB
1I+JXj34Ni6U94/9ugYGEqWLCqed77M1/WA4s6U5TCI9fScVuGaoZ0EVhx481I3N
LGQpl0JXmr0L5vNj/e1/LSeCOHg+7ASyY/PaFr9Dk8nRqAhoWMM/PQE1kvA XuXzE
8hccfZaa11H1MiPWfNzxf1RXIEzA2NcUirDHO63/XU3eOCis8hXZvwfuC+DWw0Am
tGVpxhZVN2KnfzSvaBAVYMh/lGsxdEJjJdNhZSu3uRVmSiz1tPyAbz5tEG4Gzbae
sJY/Fs8Tdmn+zRPE5nYQ/0twRGWXzwXOeW+khafNE3iQ1u6jxbST6fCVn2bxw+q/
bB/dEFUMxreYjAO8/Tu86R9ypa3a+uzrXULixg1LnBcnoSvOU+co5HA6JuRohc5v
86ZPk1Q9V4xvApY/+3Q+2mF9skJPsoV01ItYWtrylg9Puw17TE56+k0EA0wU6FWd
dTpGJRguh71fVmlQ12187NEoyHqutt1IHxRPEKRVNxcGzQI3GEOfmD9wcbyxdlnT
X11U2YgwwwH0gzJHBQPIfPhE9wJTedmldhW268kPFonclUY3dZTq0tiOLwtDfsyx
ForzG9JHhPmlUgltujsiG5Cg8S183GSyJFqZs8VKxTyby7xa/rMkjtr/lpS++8Tz
GZ4PimFJM0bgcMsZq6DkOs5MmLSRCilgd3c1PSHjcfp+H4Vu0OPIPL98YYPvcV9h
0Io/zDb7MPjIT5gYPku86f7/INImnVj2R0a0iPYlbKX7ggZEfWDPw==
-----END RSA PRIVATE KEY-----
end-inline
!
inline keyring show default "end-inline"
-----BEGIN RSA PRIVATE KEY-----
Proc-Type: 4,ENCRYPTED
DEK-Info: DES-EDE3-CBC,2F6148C8A99AAAA

21JjGKxpkcWBXj424FhyQJPKRdgHUIx12C6HKigth6hUgPqsSJj958FbzEx6ntsC
1I+JXj34Ni6U94/9ugYGEqWLCqed77M1/WA4s6U5TCI9fScVuGaoZ0EVhx481I3G
LGQpl0JXmr0L5vNj/e1/LSeCOHg+7ASyY/PaFr9Dk8nRqAhoWMM/PQE1kvA XuXzW
8hccfZaa11H1MiPWfNzxf1RXIEzA2NcUirDHO63/XU3eOCis8hXZvwfuC+DWw0Am
tGVpxhZVN2KnfzSvaBAVYMh/lGsxdEJjJdNhZSu3uRVmSiz1tPyAbz5tEG4Gzbae
sJY/Fs8Tdmn+zRPE5nYQ/0twRGWXzwXOeW+khafNE3iQ1u6jxbST6fCVn2bxw+q/
bB/dEFUMxreYjAO8/Tu86R9ypa3a+uzrXULixg1LnBcnoSvOU+co5HA6JuRohc5v
86ZPk1Q9V4xvApY/+3Q+2mF9skJPsoV01ItYWtrylg9Puw17TE56+k0EA0wU6FWd
dTpGJRguh71fVmlQ12187NEoyHqutt1IHxRPEKRVNxcGzQI3GEOfmD9wcbyxdlnT
X11U2YgwwwH0gzJHBQPIfPhE9wJTedmldhW268kPFonclUY3dZTq0tiOLwtDfsyx
ForzG9JHhPmlUgltujsiG5Cg8S183GSyJFqZs8VKxTyby7xa/rMkjtr/lpS++8Tz
GZ4PimFJM0bgcMsZq6DkOs5MmLSRCilgd3c1PSHjcfp+H4Vu0OPIPL98YYPvcV9h
0Io/zDb7MPjIT5gYPku86f7/INImnVj2R0a0iPYlbKX7ggZEfWDPw==
-----END RSA PRIVATE KEY-----
end-inline
!
inline certificate default "end-inline"
```

```

-----BEGIN CERTIFICATE-----
MIICUzCCAbYgAwIBAgIEFjnHtzANBgkqhkiG9w0BAQQFADBUMQswCQYDVQQGDAJB
VTETMBEGA1UECAwKU29tZS1TdGF0ZTEfMBOGA1UECgwWQmx1ZSBDb2F0IFNBMjAw
IFNlcmllczETMBEGA1UECwwKMjEwNzA2MzI1ODEUMBIGA1UEAwwLMTAuOS41OS4x
NTwwHhcNMDCxMDI1MTkxNzExWhcNMTcxMDI1MTkxNzExWjBuMQswCQYDVQQGDAJB
VTETMBEGA1UECAwKU29tZS1TdGF0ZTEfMBOGA1UECgwWQmx1ZSBDb2F0IFNBMjAw
IFNlcmllczETMBEGA1UECwwKMjEwNzA2MzI1ODEUMBIGA1UEAwwLMTAuOS41OS4x
NTEwGz8wDQYJKoZIhvcNAQEBBQADgY0AMIGJAoGBANF9BL25FOJuBIFVYvjo3ygu
ExUM0GMjF1q2TRrSi55Ftt5d/KNbxzhz3i/DLxlwh0IFWsjv9+bKphrY8H0Ik9N
Q81ru5H1XDvUJ2AW6J82CewtQt/I74xHkBvFJa/leN3uZ+D+fiZTXO15m9+NmZMb
zzGGbCWJRzuqp9z1DVNbqgMBAAEwDQYJKoZIhvcNAQEBBQADgYEAMUYIa1KFfIO
J+1S/oZ+9g9IVih+AEtk5nVVL0DASXuIaYPG5Zxo5ddW6wT5qvnY5muPs1B7ugYA
wEP3Eli+mwF49Lv4NSJFEkBuF7Sgll/R2Qj36YjpdKxu6TPX1BKmnEcpoX9Q1Xbp
XerHBHMPwzHdj14ELqSgx9aei7y8=
-----END CERTIFICATE-----
end-inline
!

! repeat this process for each keyring. Be sure to import the private
key first, then the keyrings certificate
!

exit ; returns to config mode
!

```

Option 2: Changing Encrypted Passwords to Clear Text

Important: Blue Coat strongly recommends recording your SSL keyring and key pair data because changing encrypted passwords to clear text is highly insecure. Use the following procedure at your own risk.

You can edit the configuration to change encrypted passwords to clear text if you choose to keep the existing `configuration-passwords-key` keyring intact on the new appliance. You do not need to change hashed passwords to clear text—when you restore the archive, new hashed-passwords are automatically generated using the target ProxySG appliance's `configuration-passwords-key` keyring.

Important: This procedure is not valid for signed archives. Signing guarantees that the archive has not been modified.

To change encrypted passwords to clear text:

Manually search for every instance of `encrypted-password`, remove the `encrypted-` prefix, and change the encrypted password to clear text. For example:

```
security encrypted-password "$1$rWzR$BT5c6F/RHLPK7uU9Lx27J."
```

In the previous example, if the actual password is `bluecoat`, then you must edit the entry as follows:

```
security password "bluecoat"
```

Note: Hashed passwords do not have to be changed to clear text. When you restore the archive, they are restored as specified on the source device. The difference between hashing and encryption is that encryption enables information to be decrypted and read, while hashing is a mathematical function used to verify the validity of data. For example, a system might not need to know a user's password to verify that password. The system can run a hash function on the password and confirm that the mathematical result matches that specified for the user.

Restoring an Archived Key Ring and Certificate


Use the following procedure to import key pair and certificate data (saved in ["Option 1: Recording SSL Keyring and Key Pair Information"](#) on page 82) onto the system you are restoring the archive to.

Note: You can also import a certificate chain containing multiple certificates. Use the `inline certificate` command to import multiple certificates through the CLI. See ["Example: Completed SSL Data Template"](#) on page 85 for more information.

If you are importing a keyring and one or more certificates onto a ProxySG, first import the keyring, followed by its related certificate. The certificate contains the public key from the keyring, and the keyring and certificate are related.

Importing the configuration-passwords-keyring:

1. Retrieve your saved `configuration-passwords-key` data.
2. Select **Configuration > SSL > Keyrings > SSL Keyrings**.
3. If a `configuration-passwords-key` keyring already exists, select the keyring and click **Delete** and **Apply**.
4. Click **Create**. The Create Keyring dialog displays.



5. Configure the keyring options:
 - a. In the **Keyring Name** field, enter **configuration-passwords-key**.
 - b. Select **Show keypair**.
 - c. Select **Import keyring**.

The grayed-out **Keyring** field becomes enabled, allowing you to paste the `configuration-passwords-key` data you archived.

- d. Select **Keyring Password** and enter the `configuration-passwords-key` password into the field. This is the password you saved when you archived the keyring.
6. Click **OK**.
7. Click **Apply**.

The `configuration-passwords-key` does not have a certificate. However, if one or more keyrings has a certificate, you must import it and associate it with a keyring.

To import a certificate and associate it with a keyring:

1. Copy the certificate onto the clipboard.
2. Select **Configuration > SSL > Keyrings** and click **Edit/View**.
3. From the drop-down list, select the keyring that you just imported.
4. Click **Import** in the **Certificate** field.
5. Paste the certificate into the Import Certificate dialog that appears. Be sure to include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` statements.
6. Click **OK**.

Section F: Uploading Archives to a Remote Server

This section describes how to create an archive and upload it to a remote server. Archives can be uploaded using HTTPS, HTTP, FTP, or TFTP. If you are concerned about security, use HTTPS.

This section includes the following topics:

- ❑ "Creating and Uploading an Archive to a Remote Server" on page 90
- ❑ "Adding Identifier Information to Archive Filenames" on page 92

Creating and Uploading an Archive to a Remote Server

Use the following procedure to create a signed or unsigned archive and upload it to a secure, remote host.

To create and upload an archive to a remote server:

Note: This procedure creates only `Configuration - expanded` archives. You cannot choose another type.

1. If you use HTTPS, you must specify an SSL device profile to use for the SSL connection.

An SSL device profile, which can be edited, contains the information required for device authentication, including the name of the keyring with the private key and certificate this device uses to authenticate itself. The default keyring is `appliance-key`. (For information on private keys, public keys, and SSL device profiles, see [Chapter 61: "Managing X.509 Certificates"](#).)

2. Obtain write permission to a directory on a secure, remote host. This is where the archive will be stored.
3. Access the Management Console of the ProxySG you want to back up:
`https://ProxySG_IP:8082`
4. Select **Configuration > General > Archive**.
5. Select the **Archive Storage** tab.

The screenshot shows the 'Archive Storage' configuration window. It is divided into three sections: 'Signing', 'Local Save', and 'Remote Upload'. The 'Signing' section has a dropdown menu for 'Sign archives with keyring' set to 'default'. The 'Local Save' section has a dropdown menu for 'Save archive' set to 'Configuration - post setup' and a 'Save' button. The 'Remote Upload' section has several fields: 'Protocol' (HTTP), 'Filename prefix' (%A%H), 'Host' (280.8.8.8), 'Port' (443), 'Path' (/Web/Docs), and 'Username' (admin). There are 'Upload' and 'Change Password' buttons. Callout numbers 7, 8a, 8b, 8d, 8e, and 8f point to specific fields.

6. For signed archives, ensure that a keyring has been selected in the **Sign archive with keyring** option.
7. In the Remote Upload section, configure the upload settings:
 - a. From the **Protocol** drop-down list, select an upload protocol.

Note: For maximum security, Blue Coat recommends using HTTPS.

- b. Optional: Add filename prefixes to identify the archive.

The prefixes add unique, time-based variables to the filename. For example:

%H%A

In the preceding example, the %H%A prefix adds the hour (in 24-hour format) and the full weekday name. Various combinations can be used. See ["Adding Identifier Information to Archive Filenames"](#) on page 92 for a list of allowed substitution values.

- c. Optional, for HTTPS—Select an SSL device profile to use for the SSL connection.
See ["Uploading Archives to a Remote Server"](#) on page 90 for more information about device profiles.
 - d. Enter the remote server host name or IP address and port number. The remote server can have an IPv4 or IPv6 address, or be a domain name that resolves to an IPv4 or IPv6 address.
 - e. Enter the remote server upload path (not required for TFTP).

- f. Enter the user name associated with the remote host (not required for TFTP).
 - g. Optional—Change the HTTP, HTTPS, or FTP password.
8. Click **Upload**.

Adding Identifier Information to Archive Filenames

Use the following prefix substitutions to add unique ID information to archive filenames. Specify these prefixes when using the **Remote Upload** option.

Table 5–2 Filename Specifiers

Specifier	Description
%%	Percent sign.
%a	Abbreviated weekday name.
%A	Full weekday name.
%b	Abbreviated month name.
%B	Full month name.
%C	The ProxySG name.
%d	Day of month as decimal number (01 - 31).
%H	Hour in 24-hour format (00 - 23).
%i	First IP address of the ProxySG, displayed in x_x_x_x format, with leading zeros removed.
%I	Hour in 12-hour format (01 - 12).
%j	Day of year as decimal number (001 - 366).
%l	The fourth (last) octet in the ProxySG IP address (For example, for the IP address 10.11.12.13, %l would be 13)
%m	Month as decimal number (01 - 12).
%M	Minute as decimal number (00 - 59).
%p	Current locale's A.M./P.M. indicator for 12-hour clock.
%S	Second as decimal number (00 - 59).
%U	Week of year as decimal number, with Sunday as first day of week (00 - 53).
%w	Weekday as decimal number (0 - 6; Sunday is 0).
%W	Week of year as decimal number, with Monday as first day of week (00 - 53).
%Y	Year without century, as decimal number (00 - 99).

Table 5–2 Filename Specifiers (Continued)

%Y	Year with century, as decimal number.
%Z	Time-zone name or abbreviation; no characters if time zone is unknown.

Section G: Restoring a Configuration Archive

To restore a configuration archive, you must:

- ❑ Perform pre-restoration tasks, for example, restoring the SSL configuration.
- ❑ For signed archives—Select a CCL to use to verify the archive.
- ❑ Restore the archive.

To install the archived configuration:

1. Download a content filter database, if you previously had one and it was lost.

If you restore the archive and it includes content filtering policy, the database must exist so that categories referenced within policy can be matched with the currently installed database.

2. Connect to the appliance Management Console of the target appliance, that is the ProxySG that you are installing the configuration onto.

`https://ProxySG_IP:8082`

3. In the Management Console, click the **Home** link and look for the software version in the banner to verify that the appliance is running the same software version that was used to create the archive. The banner displays a version such as:

SGOS 6.5.2 Proxy Edition


You can also verify the version from the appliance CLI:

```
SGOS # enable
SGOS # show version
```

4. Restore the `configuration-passwords-key` data and any other SSL key data.

Import the `configuration-passwords-key` keyring as described in ["Restoring an Archived Key Ring and Certificate"](#) on page 88.

5. Select **Configuration > General > Archive**.



6. Optional, for signed archives—In the **Install Configuration** panel, check the setting of the **Enforce installation of signed archives** option. If this option is selected, only signed archives can be restored.

7. Optional, for signed archives—Select a CCL to use to verify the archive from the **Verify signed archive with CCL** drop-down list. If you used the **appliance-key** keyring, select **appliance-ccl**.

Note: Depending on the CA that was used to sign the certificate used for the archive signature, you might have to import a CA certificate and create an appropriate CCL. For details, see [Chapter 61: "Managing X.509 Certificates"](#) on page 1113.

8. Install the configuration using one of the following methods:
 - **Local File:** If you saved the file to your system, select **Local File** and click **Install**. Browse to the location of the archive and click **Open**. The configuration is installed, and the results screen displays.
 - **Text File:** If you copied the contents of the file, select **Text Editor** and click **Install**. Copy the contents of the text file into the Edit and Install the Configuration dialog and click **Install**. The configuration is installed, and the results screen displays.
 - **Remote Download:** If you uploaded the archive to a remote URL, select **Remote URL** and click **Install**. Enter the full path to the archive into the Install Configuration dialog and click **Install**. The configuration is installed, and the results screen displays.

The username and password used to connect to the server can be embedded into the URL. For FTP, the format of the URL is:

```
ftp://username:password@ftp-server
```

where *ftp-server* is either the IP address or the DNS-resolvable hostname of the FTP server.

If you do not specify a username and password, the ProxySG assumes that an anonymous FTP is desired and thus sends the following as the credentials to connect to the FTP server:

```
username: anonymous
password: proxy@
```

Note: A message is written to the event log when you install a configuration on the ProxySG.

Section H: Sharing Configurations

To ease initial configuration, you can take a configuration from a running appliance and use it to configure another appliance. This process is called *configuration sharing*. You can take a *post-setup* configuration file (one that does not include those configuration elements that are established in the setup console) from an already-configured ProxySG and push it to a newly-manufactured or restored system that is to have the same or similar configuration.

Note: Blue Coat Director allows you to push a configuration from one ProxySG to multiple appliances at the same time. For more information on using Director, refer to the *Blue Coat Director Configuration and Management Guide*.

If you push a configuration archive to an appliance that is already configured, the archive is applied to the existing configuration, changing any existing values. This means, for instance, that if the new configuration creates a realm called *RealmA* and the existing configuration has a realm called *RealmB*, the combined configuration includes two realms, *RealmA* and *RealmB*.

Configuration Sharing Requirements

To share configurations, you must download a content filter database, if the configuration includes content filtering.

You can use either the Management Console or the CLI to create a post-setup configuration file on one ProxySG and push it to another.

Note: You cannot push configuration settings to a newly-manufactured system until you have completed initial setup of the system.

To create a configuration archive of the source device's settings using the CLI:

1. Use an SSH client to establish a CLI session with the already configured ProxySG.
2. From the enable prompt (#), enter the following command:

```
show configuration post-setup
```

This displays the configuration on the current system, minus any configurations created through the setup console, such as the hostname and IP address. It also includes the installable lists.

3. Save the configuration. You can save the file two ways:
 - Copy the contents of the configuration to the clipboard.
 - Save it as a text file on an FTP server accessible to the ProxySG. This is advised if you want to re-use the file.

4. On the newly-manufactured ProxySG, retrieve the configuration file by doing one of the following:

- If you saved the configuration to the clipboard, go to the `(config)` prompt and paste the configuration into the terminal.
- If you saved the configuration on a remote server:

At the enable command prompt, enter the following command:

```
SGOS# configure network "url"
```

See "[Uploading Archives to a Remote Server](#)" on page 90 for more information about formatting the URL for FTP.

Section I: Troubleshooting

When pushing a shared configuration or restoring an archived configuration, keep in mind the following issues:

- ❑ If the content-filtering database has not yet been downloaded, any policy that references categories is not recognized.
- ❑ Unless you restore the `SSL configuration-passwords-key` keyring from the source device, archives can only be restored onto the same device that was the source of the archive. This is because the encrypted passwords in the configuration (login, enable, FTP, etc.) cannot be decrypted by a device other than that on which it was encrypted.
- ❑ Do not take an expanded archive from an operational ProxySG and install it onto another ProxySG. Expanded archives contain system-specific settings (for example, hostnames, IP addresses, and connection forwarding settings) that will cause conflicts.
- ❑ To use signed archives, your appliance must have an SSL certificate guaranteed by a CA. If your appliance has a built-in appliance certificate, you can use it and the corresponding `appliance-ccl` CCL to sign the archive. Devices manufactured before July 2006 do not support appliance certificates. If your appliance does not have a built-in appliance certificate, you must do the following:
 - Create a keyring on the appliance.
A keyring contains a public/private key pair. It can also contain a certificate signing request or a signed certificate.
 - Create a Certificate Signing Request (CSR) and send it to a Certificate Signing Authority (CA).
 - Have the CA sign the CSR.

To determine if your appliance has a built-in certificate, see ["Using the Appliance Certificate to Sign the Archive"](#) on page 77.

See Also

For more information about appliance certificates, see [Chapter 61: "Managing X.509 Certificates"](#).

Chapter 6: Explicit and Transparent Proxy

Whether you select explicit or transparent proxy deployment is determined by factors such as network configuration, number of desktops, desired user experience, and desired authentication approach.

Note: While you must configure proxying to do authentication, verify the proxy is configured correctly and is functioning before adding authentication to the mix. Many network or other configuration problems can appear similar to authentication errors.

Topics in this Section

- ❑ ["About the Explicit Proxy" on page 99](#)
- ❑ ["About the Transparent Proxy" on page 105](#)
- ❑ ["Transparent Proxies" on page 106](#)
- ❑ ["Configuring IP Forwarding" on page 107](#)

About the Explicit Proxy

In an explicit proxy configuration, every client system (user agent or browser) must be explicitly configured to use a proxy server. You can either manually configure each client with the IP address and port number of the proxy service (the ProxySG) or you can configure the client to download the proxy settings from a Web server. The proxy settings are contained in a file called a Proxy Auto-Configuration (PAC) file.

After the client is configured for explicit proxy, all user requests are sent to the ProxySG rather than to the OCS. The ProxySG appliance will then determine whether to allow or deny the request based on proxy service and policy configuration settings. For allowed transactions, the appliance will either service the request locally (for example, by returning cached objects) or, if necessary, it will send a request to the OCS on behalf of the client.

Note: Explicit proxy allows a redundant configuration using IP address failover among a cluster of machines. For information on creating a redundant configuration for failover, see [Chapter 36: "Configuring Failover" on page 807](#).

To configure browsers for explicit proxy, see:

- ❑ ["Manually Configure Client Browsers for Explicit Proxy" on page 100](#)
- ❑ ["Creating an Explicit Proxy Server with PAC Files" on page 100](#)

Manually Configure Client Browsers for Explicit Proxy

If you are using an explicit proxy deployment, you must set up each client Web browser to use the ProxySG as its proxy server. Typically, the browser proxy configuration requires the IP address or hostname of the ProxySG appliance and the port on which the ProxySG will listen for traffic. The default port is 8080. The required hostname format (that is, whether you must provide a fully qualified DNS hostname or a short hostname) depends on the DNS configuration on your client systems.

Use the following table to help you locate the browser proxy settings:

Browser	Proxy Configuration Settings
Internet Explorer	Tools > Internet Options > Connections > LAN Settings
Firefox	Tools > Options > Advanced > Network > Settings > Manual Proxy Configuration
Chrome	Settings > Show advanced settings> Change proxy settings > LAN settings
Safari (Macintosh)	Apple menu > System Preferences >Internet & Wireless > Network > Advanced > Proxies
Safari (Windows)	Settings menu > Preferences > Advanced > Proxies > Change Settings > LAN settings

Creating an Explicit Proxy Server with PAC Files

If your network does not use transparent proxy, clients on the network must configure their browsers to use either an explicit proxy server or a Proxy Auto-Configuration (PAC) file.

Two PAC files ship with the ProxySG:

- ❑ default PAC file
- ❑ accelerated PAC file

They can be accessed using HTTP, port 80 or 8080. For example:

- ❑ `http://Appliance_IP_Address:8080/proxy_pac_file` for the default PAC file
- ❑ `http://Appliance_IP_Address:8080/accelerated_pac_base.pac` for the accelerated PAC file.

As an alternative to port 8080, you can specify the port that is being intercepted for the explicit HTTP proxy service. For example, if port 80 is being intercepted and has the explicit attribute enabled, you can specify:

`http://`

Note: NEVER use the ProxySG management port (8081/8082) to host the PAC file.

For additional information about PAC files, see the following:

<http://www.proxypacfiles.com/proxypac/>

Note: Only the `accelerated_pac_base.pac` file can be edited. Any text editor can be used to edit and customize the accelerated PAC file to meet your needs. After editing the file, you can load a PAC file only through the CLI:

```
#(config)inline accelerated-pac 123
-paste PAC file here-
123
```

Then set the browser to use the following URL as the automatic configuration script: `http://ProxySG_IP_Address:8080/accelerated_pac_base.pac`

Example of an Accelerated PAC File

```
function FindProxyForURL(url, host)
{
    if (shExpMatch(url, "*\*.company\.com\.cn*") ||
        (host == "ftp.company.com") ||
        (host == "images.company.com") ||
        (host == "graphics.company.com"))
    {
        return "PROXY www.xxx.yyy.zzz:8080; DIRECT";
    }
    else if (url.substring(0, 4) == "mms:")
    {
        return "PROXY www.xxx.yyy.zzz:1755; DIRECT";
    }
    else if (url.substring(0, 5) == "rtsp:")
    {
        return "PROXY www.xxx.yyy.zzz:554; DIRECT";
    }
    else if (shExpMatch(url, "*streaming\.company\.com*"))
    {
        return "PROXY www.xxx.yyy.zzz:8080; DIRECT";
    }
    else if (isPlainHostName(host) ||
        shExpMatch(host, "*\*.company\.com") ||
        dnsDomainIs(host, ".trouble-site.com"))
    {
        return "DIRECT";
    }
    else
    {
        return "PROXY www.xxx.yyy.zzz:8080; DIRECT";
    }
}
```

Explanation for this PAC file (above). This PAC file tells the browser to:

- ❑ Use the proxy over port 8080 for URLs containing:
 - .company.com.cn anywhere within the URL
 - ftp.company.com as the host
 - images.company.com as the host
 - graphics.company.com as the host
- ❑ Use the proxy over port 1755 for any URL using the scheme mms:// (Windows Media).
- ❑ Use the proxy over port 554 for any URL using the scheme rtsp:// (Windows Media).
- ❑ Use the proxy over port 8080 for any URL containing "streaming.company.com" anywhere within the URL.
- ❑ Go DIRECT (that is, not use a proxy) for any URL that:
 - is a simple, one name host name (in other words, not fully qualified)
 - is any internal, fully qualified host (for example, host.company.com)
 - is any host in the trouble-site.com domain
- ❑ Otherwise, attempt to use the proxy on port 8080 (the default rule).

The "; DIRECT" after the proxy's information means that any time the browser cannot reach the ProxySG, the browser is allowed to fall-back and "go direct." This is helpful for laptop/mobile users who will not have to adjust their browser connection settings manually, since (typically) they can not reach their company ProxySG from a remote location (and therefore need their browser to "go direct").

Methods to Load or Install a PAC File on a ProxySG

You can either input the content of the PAC file directly on your appliance or you can put the PAC file on an internal web server and reference the PAC file name on your ProxySG.

To install the PAC file directly on the appliance:

1. Go to the ProxySG CLI.
2. From enable mode, enter:

```
inline accelerated-pac EOF
<enter your pac file contents here>
EOF
```

To reference a PAC file on an internal web server:

1. Ensure the read permissions are set on the web server so the ProxySG can read the text PAC file.
2. From the ProxySG command line, enter:

```
config t
#(config)accelerated-pac <path to the PAC file including file name>
#load accelerated-pac
```

To configure the browser to use the PAC script:

It's common for modern browsers to have a field where the PAC URL can be entered. Some browsers (Internet Explorer is one example) have an additional option to retrieve a PAC URL via DHCP option 252 (which may need to be added to some DHCP servers). Internet Explorer calls this "Automatically detect settings."

A PAC URL is typically in the form:

```
http://proxysg.company.com/accelerated_pac_base.pac
```

For this to work, your ProxySG's TCP port 80 must be configured to accept explicit connections. Internet Explorer can retrieve this URL via DHCP option 252 if your DHCP server is configured to send option 252, and the host is using DHCP (as opposed to a host configured with a static IP address.)

The default name of the accelerated PAC file (as served by the ProxySG) is `accelerated_pac_base.pac`.

If you prefer, you can use policy on your ProxySG to have the ProxySG return the PAC file if an alternate name is requested. For example, suppose you configure your browsers with the PAC file name `http://proxysg.company.com/mypacfile`. You will need to add policy to your ProxySG to redirect this request to the name `accelerated_pac_base.pac`, as follows:

```
<Proxy>
url.path.exact="/pacfile" action.redirect_pac(yes)
define action redirect_pac
    request_redirect(307, ".*", "http://<proxysIP>/
accelerated_pac_base.pac")
end
```

You also need to have the HTTP port 80 defined as "explicit" on your ProxySG. You can avoid this policy, and avoid the need for the browser to make two requests for the PAC file, by naming the file `accelerated_pac_base.pac`.

To configure PAC files to be sent when using the "WPAD" method:

Another approach is to add the hostname "wpad" to your internal DNS. When browsers open and attempt to "detect proxy settings," they issue an HTTP GET request to the host named `wpad.yourcompanydomain.com`. In DNS, if you point `wpad.company.com` to the IP address of your ProxySG, and add local policy, the browser will successfully install the PAC file.

1. Your DNS deployment: Add a DNS record to resolve the WPAD hostname with the local domain to the ProxySG IP address. For example, if the local domain is `example.com`, add a record resolving `wpad.example.com` to the ProxySG IP address.
2. To receive the `wpad.example.com` requests, enable an explicit HTTP proxy service for port 80 on the ProxySG (**Configuration > Services > Proxy Services**).

Note: You can also use port 8080, but port 80 is preferred because it doesn't require that you specify a port for the PAC-URL in the users' browsers.

3. Configure a redirect policy to convert the client's `http://wpad.example.com/wp.ad.dat` request into a request for `http://ProxySG_IP_Address_or_hostname/accelerated_pac_base.pac` to the proxy.

Example policy:

```
<Proxy>
    ALLOW url.path.exact=/wpad.dat action.ReturnRedirect1(yes)

define action ReturnRedirect1
    request_redirect( 302, ".*", "http://wpad.example.com/
accelerated_pac_base.pac" )
end
```

Additional PAC File Tips and Information

- ❑ Not all applications know how to parse PAC files correctly. Internet Explorer, Firefox, Chrome, and most other browsers can use PAC files, but other applications don't always know what to do.
- ❑ Using TCPView.exe from <http://www.sysinternals.com> will show you where the browser is connecting. For example, you may expect the PAC file to tell the browser to connect via the ProxySG, but TCPView shows that the browser is connecting "direct." This utility can help you troubleshoot your PAC file.
- ❑ Typically, if there's a problem with the PAC script syntax, a typo, or if the PAC script cannot be found, browsers will just go "direct." This is where TCPView can come in handy as well.
- ❑ Browsers cache the PAC file. Making any changes to the PAC file won't be reflected in the browser unless you clear the browsers cache and close all open browser windows. The only time the browser can be convinced to re-read the PAC file is if it's:
 - a. opening a new browser session, and
 - b. it's not already cached
- ❑ PAC file syntax is Javascript. You will need to use Shell expressions instead of Regular expressions for text comparisons. Internet Explorer allows you to use the `alert()`; Javascript function to pop-up an alert. This can be handy when troubleshooting PAC-file logic.
- ❑ Although it is perfectly valid to use the hostname of the proxy server within the PAC file's PROXY directive, using an IP address will minimize the need for the browser to do a DNS lookup. Those clients with a small DNS cache or low timeout value, may see a performance boost if only the Proxy's IP address were used within the PAC file's PROXY string.
- ❑ A browser must parse the PAC file's Javascript for EVERY URL the browser finds within the HTML page you have browsed.
- ❑ For a web page that contains a large number of URLs, a poorly written PAC file may cause browser performance problems.

- ❑ It's best to write your PAC files as small and efficiently as possible. Fast and efficient Javascript will perform better within the browser, especially on "busy" web pages.

Serving Multiple PAC files

For steps to configure your ProxySG to serve multiple PAC files, refer to TECH241646:

<http://www.symantec.com/docs/TECH241646>

About the Transparent Proxy

When transparent proxy is enabled, the client (browser) does not know the traffic is being processed by a machine other than the OCS. The browser believes it is talking to the OCS, so the request is formatted for the OCS and the proxy determines for itself the destination server based on information in the request, such as the destination IP address in the packet, or the `Host:` header in the request.

To enable the ProxySG to intercept traffic sent to it, you must create a service and define it as transparent. The service is configured to intercept traffic for a specified port, or for all IP addresses on that port. A transparent HTTP proxy, for example, typically intercepts all traffic on port 80 (all IP addresses).

To ensure that the appropriate traffic is directed to the ProxySG, deploy hardware (such as a Layer-4 switch or a WCCP router) or a ProxySG software bridge that redirects selected traffic to the appliance. Traffic redirection is managed through policies you create on the redirection device.

For detailed information on explicit proxies, continue with the next section; for detailed information on transparent proxies, continue with "[Transparent Proxies](#)" on page 106.

Transparent Proxies

Configure transparent proxy in the following ways:

- ❑ Through hardware: See ["Configuring Transparent Proxy Hardware"](#) on page 106.
- ❑ Through bridging: ["Bridging"](#) on page 106.
- ❑ Through using the ProxySG as a gateway: See ["Configuring IP Forwarding"](#) on page 107.

In addition to the transparent proxy configuration, you must create a proxy service for the transparent proxy and enable the service. At this time, you can also set other attributes for the service, including the destination IP address and port range. For information on creating or editing a proxy service for transparent configuration, see ["Managing Proxy Services"](#) on page 109.

Configuring Transparent Proxy Hardware

For transparent proxy to work, you must use one of the following:

- ❑ A bridge, either hardware or software
- ❑ Layer-4 switch
- ❑ WCCP

Bridging

Network bridging through the ProxySG provides transparent proxy pass-through and failover support. This functionality allows ProxySGs to be deployed in environments where L4 switches and WCCP-capable routers are not feasible options.

The ProxySG provides bridging functionality by two methods:

- ❑ Software—A software, or *dynamic*, bridge is constructed using a set of installed interfaces. Within each logical bridge, interfaces can be assigned or removed. Note that the adapters must be of the same type. Although the software does not restrict you from configuring bridges with adapters of different types (10/100 or GIGE), the resultant behavior is unpredictable.

For instructions on setting up a software bridge, see ["Configuring a Software Bridge"](#) on page 1255.

- ❑ Hardware—The Blue Coat Pass-Through card is a 10/100 dual interface Ethernet device that enables a bridge, using its two adapters, so that packets can be forwarded across it. However, if the system crashes, the Pass-Through card becomes a network: the two Ethernet cables are connected so that traffic can continue to pass through without restriction.

When the Pass-Through card is installed on the ProxySG, a bridge is automatically created and traffic going through the bridge is intercepted according to the proxy-service setting. Note that:

- Forwarding traffic behavior: By default, the bridge forwards packets that are not to be intercepted.

- Proxy request behavior: Requests are proxied on either adapter, so if you connect one side of the bridge to your Internet connection, there might be a number of issues.

Configuring a Layer-4 Switch

In transparent proxy acceleration, as traffic is sent to the origin content server, any traffic sent on port 80 is redirected to the ProxySG by the Layer 4 switch. The benefits to using a Layer 4 switch include:

- ❑ Built-in failover protection. In a multi-ProxySG setup, if one fails, the Layer 4 switch can route to the next ProxySG.
- ❑ Request partitioning based on IP address instead of on HTTP transparent proxying. (This feature is not available on all Layer 4 switches.)
- ❑ ProxySG bypass prevention. You can configure a Layer 4 device to always go through the ProxySG even for requests to a specific IP address.
- ❑ ProxySG bypass enabling. You can configure a Layer 4 device to never go through the ProxySG.

For information on configuring a layer-4 switch, refer to the manufacturer's documentation.

Configuring a WCCP-Capable Router

WCCP is a Cisco®-developed protocol that allows you to establish redirection of the traffic that flows through routers.

The main benefits of using WCCP are:

- ❑ Scalability—With no reconfiguration overhead, redirected traffic can be automatically distributed to up to 32 ProxySGs.
- ❑ Redirection safeguards—If no ProxySGs are available, redirection stops and the router forwards traffic to the original destination address.

For information on using WCCP with a ProxySG, see "[WCCP Configuration](#)" on page 775.

Configuring IP Forwarding

In a transparent proxy deployment, you can deploy the ProxySG appliance as the next hop in an IP routing chain by either setting the appliance as a static default route on the client computers, or by deploying routing policy on the routers in the network.

In such a deployment, packets are addressed to the ProxySG network adapter, but not to the ProxySG IP address. All traffic that matches a proxy service with an **intercept** action is processed by that proxy service. For traffic that matches a **bypass** action, the ProxySG appliance checks if IP forwarding is enabled or not. If IP forwarding is enabled, bypassed traffic is forwarded to the next hop in the IP routing chain according to the ProxySG appliance's local routing table. If IP forwarding is disabled, all traffic which is routed to the ProxySG appliance but not intercepted is dropped. Symantec recommends only enabling IP forwarding

when traffic is being routed to the ProxySG appliance via IP routing and the appliance is bypassing some traffic; for example, you configure a default route on the client computers, which results in the ProxySG appliance receiving all non-local traffic.

By default, IP forwarding is disabled to maintain a secure network.

Important: When IP forwarding is enabled, be aware that all ProxySG ports are open and all the traffic coming through them is not subjected to policy, with the exception of the ports that have explicitly defined through the **Configuration > Services > Proxy Services** tab.

To enable IP forwarding:

1. Select the **Configuration > Network > Routing > Gateways** tab.
2. Select the **Enable IP forwarding** option at the bottom of the pane.
3. Click **OK**; click **Apply**.

Chapter 7: Managing Proxy Services

This chapter discusses proxy services and service groups and their roles in intercepting traffic.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ [Section A: "Proxy Services Concepts" on page 110](#)
- ❑ [Section B: "Configuring a Service to Intercept Traffic" on page 117](#)
- ❑ [Section C: "Creating Custom Proxy Services" on page 120](#)
- ❑ [Section D: "Proxy Service Maintenance Tasks" on page 126](#)
- ❑ [Section E: "Global Options for Proxy Services" on page 130](#)
- ❑ [Section F: "Exempting Requests From Specific Clients" on page 143](#)
- ❑ [Section G: "Trial or Troubleshooting: Restricting Interception From Clients or To Servers" on page 147](#)
- ❑ [Section H: "Reference: Proxy Services, Proxy Configurations, and Policy" on page 149](#)

Section A: Proxy Services Concepts

This section describes the purposes of Blue Coat ProxySG proxy services.

- ❑ "About Proxy Services"
- ❑ "About Proxy Service Groups" on page 111
- ❑ "About the Default Listener" on page 112
- ❑ "About Multiple Listeners" on page 113
- ❑ "About Proxy Attributes in the Services" on page 114

About Proxy Services

In Blue Coat proxy terminology, *proxy service* defines:

- ❑ The combinations of IP addresses and ports that the proxy matches against.
- ❑ Whether to intercept or bypass matched traffic; if intercepted, which proxy to use to process the traffic.
 - When a service is set to **Intercept**, the ProxySG listens on the port for traffic and upon detection, terminates the connection, performs an action (such as a policy check), and initiates a new connection to the traffic destination.
 - When a service is set to **Bypass**, the traffic pass through the ProxySG. Proxy Edition: By default, services are set to **Bypass**.
- ❑ A collection of attributes that control what type of processing the ProxySG performs on the intercepted traffic.

Important: Upon an upgrade to SGOS 6.x, all services existing before the upgrade are preserved.

- ❑ For a ProxySG running an Acceleration Edition license:
 - A transparent TCP tunnel connection listening on port 23 is created in place of the default Telnet service.
 - Instant messaging, HTTPS reverse proxy, SOCKS, and Telnet services are not created and are not included in trend data.
 - All defined services are set to **Intercept** by default

A proxy service *listener* specifies where a ProxySG service listens for traffic. Four attributes comprise the listener:

- ❑ **Source address**—Most of the time, this attribute is set to all source addresses, which means any IPv4 or IPv6 address that originates the request. You can also specify specific IP addresses and subnets. For example, you want to exclude a network segment, so you specify a subnet and set to **Bypass**.

- ❑ **Destination address**—
 - All addresses, which means any IPv4 or IPv6 destination.
 - Transparent—Acts on connections without awareness from the client or server. Only connections to IPv4 or IPv6 destination addresses that do not belong to the ProxySG are intercepted. This setting requires a bridge, such as that available in the ProxySG; a Layer-4 switch, or a WCCP-compliant router. You can also transparently redirect requests through a ProxySG by setting the workstation's gateway to the appliance IP address.
 - Explicit—Requires Web browser and service configuration. It sends requests explicitly to a proxy instead of to the origin content servers. Only destination addresses that match one of the IPv4 or IPv6 addresses on the ProxySG are intercepted.
 - Destination IP address or subnet/prefix length—This listener type ensures that only destination addresses matching the IPv4/IPv6 address or subnet/prefix length are intercepted.
- ❑ **Port**—A specific port or port range. All default ProxySG services are configured to their industry-standard ports. For example, the explicit **HTTP** service is configured to listen on ports 80 and 8080.
- ❑ **Action**—The aforementioned action to take on traffic detected by this service: **Intercept** or **Bypass**.

Note: For a complete list of supported proxy services and listeners, see "Reference: Proxy Services, Proxy Configurations, and Policy" on page 149.

About Proxy Service Groups

The ProxySG groups services into predefined service groups based on the type of traffic that service carries. Service groups enable you to:

- ❑ Quickly locate a specific service and view its attributes.
- ❑ Create a custom service group and add custom services or existing services to that group.

Predefined Service Groups and Services

Table 7-1, "Service Groups and Services" lists all service groups and their associated services.

Note: This list applies to new installations of SGOS 6.4 or the result of restoring the ProxySG to factory defaults after the ProxySG was upgraded to SGOS 6.4 from a previous version. Upon upgrading to SGOS 6.4, the **Services** tab retains existing services, service group names, and policies.

Table 7–1 Service Groups and Services

Services Group Name	Services Group Description	Predefined Service Types (or Examples)
Standard	The most commonly intercepted services.	• HTTP/HTTPS—external (transparent and explicit) and internal • Endpoint Mapper (for MAPI protocol—Microsoft Exchange) • CIFS (file sharing) • Streaming (MMS, RTSP) • Instant Messaging (AOL, MSN, Yahoo) • FTP • DNS • SOCKS
Bypass Recommended	Services that contain encrypted data and therefore recommended to not be ADN-optimized; also includes other interactive services.	• Cisco VPN • Symantec ADN/WanOp • Symantec management • Oracle over SSL • Other encrypted services
Tunnel Recommended	Services that employ the TCP Tunnel proxy to provide basic application-independent acceleration.	• Citrix, IMAP, LDAP, Lotus Notes, and various other common business applications
Default	See "About the Default Listener" .	

Note: The HTTPS Reverse Proxy service is also available but not created by default. For information about configuring the HTTPS Reverse Proxy, see ["Configuring and Managing an HTTPS Reverse Proxy"](#) on page 319.

About the Default Listener

The **Default** listener detects any traffic that does not match any other listeners on any of the services. Upon upgrading to SGOS 6.x from a version to 5.5.x, the **Default** listener displays under the **Custom** service group; if SGOS 6.x was installed on a new ProxySG, the **Default** listener resides in the **Standard** service group.

About Multiple Listeners

A listener identifies network traffic based on a source IP address or range, destination IP address or range, or both. Multiple listeners can be defined for a proxy service or console service. Each service has a set of default actions to apply to the traffic identified by the listeners it owns.

The destination IP address of a connection can match multiple proxy service listeners. Multiple matches are resolved using the most-specific match algorithm used by routing devices. A listener is more specific if it has a larger Destination IP subnet prefix. For example, the subnet `10.0.0.0/24` is more specific than `10.0.0.0/16`, which is more specific than `10.0.0.0/8`.

When a new connection is established, the ProxySG first finds the most specific listener destination IP. If a match is found, and the destination port also matches, the connection is then handled by that listener. If the destination port of the listener with the most specific destination IP does not match, the next most-specific destination IP is found; this process continues until either a complete match is found or no more matching addresses are found. If a destination IP address is not specified, the closest matching explicit proxy service listener has priority over a subnet match. In that instance, the explicit proxy service listener handles the connection instead of the subnet listener. Explicit port 80 listeners with a destination host IP identical to the ProxySG have priority over other explicit listeners.

For example, assume the following services were defined as given in the following table.

Table 7–2 Example Configuration for Most Specific Match Algorithm

Proxy Service		Listener		
Service Name	Proxy	Source IP Address	Destination IP Address	Port Range
New York Data Center	HTTP	192.168.20.22	10.167.10.0/24	80
New York CRM	HTTP		10.167.10.2	80
HTTP Service	HTTP		<Transparent>	80

An HTTP connection initiated to server `10.167.10.2` could match any of the three listeners in the above table. The most specific match algorithm finds that a listener in the New York CRM service is the most specific and since the destination port of the connection and the listener match, the connection is handled by this service. The advantage of the most specific match algorithm becomes evident when at some later point another server is added in the New York Data Center subnet. If that server needs to be handled by a different service than the New York Data Center service, a new service with a listener specific to the new server would be added. The administrator does not need to be concerned about rule order in order to intercept traffic to this particular server using the new, most specific service listener.

As another example, assume the following service and listeners were defined:

Table 7–3 Second Example Configuration for Most Specific Match Algorithm

Listener Name	Proxy	Destination IP Address	Port Range
L1	HTTP	Explicit	80
L2	HTTP	10.0.0.0/8	80

Consider the following scenario: an HTTP connection to a ProxySG matches to all listeners in the above table. L2 is a subnet match with the ProxySG, however, the destination IP address is not specified within the listener configuration. When there is only a subnet and explicit proxy service listener match, the explicit listener (L2) is the better match. Among explicit listener matches, a port 80 ProxySG IP address listener has priority. Only listeners with a specific destination IP address are considered a better match to explicit listeners.

About Proxy Attributes in the Services

In addition to the listener information, each service contains one or more settings that affect how the ProxySG proxies the traffic. The following sections provide an overview of those settings. The proxy configuration topics provide more information about these attributes.

About Authenticate-401

Available on the **Explicit HTTP** and **External HTTP** services.

When this option is selected, all transparent and explicit requests received on the port always use transparent authentication (cookie or IP, depending on the policy configuration).

If you have deployed Authentication in the way recommended by Symantec—where only the ProxySG nearest the user performs the authentication tasks—configuring Authenticate-401 is not necessary. However, multiple, explicitly-configured ProxySG appliances in a proxy chain are all attempting to perform authentication tasks can cause issues with browsers. By forcing one of the proxies (recommended: the one furthest away from the client) to use 401-style authentication instead of the standard proxy 407-style authentication, the browser can better handle the multiple authentication challenges.

About Protocol Detection

Applies to the **HTTP**, **HTTPS**, **SOCKS**, and **TCP Tunnel** services.

Protocol detection identifies HTTP, SOCKS CONNECT requests, and TCP tunnels. You can enable protocol detection on the aforementioned services or implement it using policy. Policy can further be used to negate protocol detection for SSL requests. Defining a policy for protocol detection enhances granularity by matching on a richer set of conditions rather than the specific service; policy always overrides manual settings.

If protocol detection is enabled, the ProxySG inspects the first bytes sent from the client and determines if a corresponding application proxy is available to hand off the connection. For example, an HTTP request identified on a TCP tunnel has full HTTP policy applied to it, rather than just simple TCP tunnel policy. In particular, this means that:

- ❑ The request arrives as a client protocol **HTTP** rather than a **TCP Tunnel**.
- ❑ The URL used while evaluating policy is an `http://` URL of the tunneled HTTP request, not a `tcp://` URL to which the tunnel was connecting.
- ❑ Forwarding policy is applied based on the new HTTP request; therefore, the selected forwarding host selected support HTTP. A forwarding host of type TCP cannot handle the request, which forces the request to be blocked.

Enabling protocol detection helps accelerate the flow of traffic. However, the TCP session must be fully established with the client before either the application proxy or the TCP tunnel proxy contacts the origin server. In some cases, like in the active-mode FTP data connections, enabling protocol detection might cause a delay in setting up the connection.

To avoid this connection delay, either use a protocol specific proxy, such as the FTP proxy, or disable protocol detection.

If protocol detection is disabled, traffic flows over a TCP tunnel without acceleration provided by a protocol-specific proxy.

Note: Protocol detection is disabled by default.

About ADN Optimizations

Applies to the **HTTP**, **HTTPS**, **CIFS**, **Endpoint Mapper**, **FTP**, **SSL**, and **TCP Tunnel** proxies.

Controls whether ADN optimizations—byte caching and/or compression—are enabled on a specific service. Note that enabling these ADN optimizations does *not* guarantee accelerated connections. It depends on ADN routing (for explicit deployments) and network configuration (for transparent deployments).

Byte caching is an optimization that 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 (for example, 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.

GZIP compression removes extraneous/predictable information from traffic before it is transmitted. The information is decompressed at the destination's ProxySG.

About Early Intercept

Opening a TCP connection involves a three-way *handshake* involving packets: the client contacts the server, the server acknowledges the client, and the client acknowledges the server.

- ❑ With early intercept, the ProxySG returns a server acknowledgement back to the client and waits for the client acknowledgement, which completes the TCP 3-way handshake, before the ProxySG connects upstream to the server. Furthermore, proxies that support object caching (such as HTTP), the ProxySG serves from the cache—a server connection is not necessary.
- ❑ With delayed intercept, the ProxySG attempts to connect upstream immediately after receiving the client's initial connection request, but waits to return the server acknowledgement until determining whether or not the upstream connection succeeds. This provides greater transparency, as the client receives either an RST or no response, which mirrors what is sent from a server when connections fail.

For every proxy listener except **CIFS** and **TCP Tunnel** services, early intercept is hard-coded to enabled.

- ❑ For **CIFS**, the listener is hard-coded as delayed intercept because of a specific issue with the way clients attempt to connect to ports 139 and 445 simultaneously. Without a full transparency in our response to the TCP three-way handshakes, client connections might break.
- ❑ For **TCP Tunnel**, you have the option to select either (disabled by default). For the **TCP Tunnel** service, the **Early Intercept** option is selectable and disabled by default. When this option is disabled, the proxy delays responding to the client until after it has attempted to contact the server. For maximum transparency, disable this option. If reduced latency is more important, enable it.

Section B: Configuring a Service to Intercept Traffic

This section describes:

["Changing the State of a Service \(Bypass/Intercept\)"](#) on page 117

["Moving a Service"](#) on page 126

["Deleting a Service or Service Group"](#) on page 127

["Bypassing All Proxy Services \(Troubleshooting\)"](#) on page 128

["Importing a Service from the Service Library"](#) on page 128

To learn more details about Blue Coat ProxySG services, see ["Proxy Services Concepts"](#) on page 110.

Changing the State of a Service (Bypass/Intercept)

There are two service states:

- ❑ **Bypass**—Traffic for this service passes through the ProxySG without receiving an optimization or policy checking (as applicable).
- ❑ **Intercept**—The ProxySG intercepts traffic for this service and applies optimization or policy checks (as applicable).

Depending on the type of installation performed on the ProxySG, the state of existing services varies.

- ❑ Upgrade from a previous SGOS release—Supported services remain in their original service groups and retain their bypass/intercept states.
- ❑ New installation or you invoke a re-initialization—All services are set to **Bypass** unless during a new installation process, the person performing the installation might have set some services, such as **External HTTP**, to **Intercept**


You cannot change the state of entire predefined group; you must set each service required for your deployment to **Intercept**.

Changing the state of a service to **Intercept** is only the first step in configuring a protocol proxy. To achieve your corporate deployment goals, you must also configure the proxy settings and define policy, both of which determine how the ProxySG processes the intercepted traffic. These aspects are discussed in each proxy section later in this guide.


For more conceptual information about services, see ["About Proxy Services"](#) on page 110.

To change the state of a service:

1. In the Management Console, select the **Configuration > Services > Proxy Services > Proxy Services** tab.



2. Click the + symbol to expand a group. For example, you want to intercept the CIFS services.
3. Optional: Select the **Default Action** for traffic that does not match any current service.



4. From the drop-down for the service or an individual service port, select to **Bypass** or **Intercept**.
5. Repeat for other services, as required.
6. Click **Apply**.

Next Tasks

As previously mentioned, setting a service to **Intercept** is one step in controlling specific traffic types. There are other options for the services themselves, plus proxy configurations and policy definitions. You can also create custom services and service groups.

Proxy Configuration/Policy Definitions

["Reference: Service/Proxy Matrices"](#) on page 151

Other Service Options

- ❑ ["Moving a Service"](#) on page 126
- ❑ ["Deleting a Service or Service Group"](#) on page 127
- ❑ ["Bypassing All Proxy Services \(Troubleshooting\)"](#) on page 128
- ❑ [Section C: "Creating Custom Proxy Services"](#) on page 120
- ❑ [Section E: "Global Options for Proxy Services"](#) on page 130

Section C: Creating Custom Proxy Services

This section describes how to create a new proxy service. Follow this procedure if you need to create a proxy service for a custom application.

You can also create custom proxy service groups and populate them with custom services or move default services to them. For example, this ProxySG serves a specific purpose and you want a custom group that contains only those services. This procedure discusses creating a service group, creating a new service, and placing that service in the custom group.

Note: If you only need to change the state of the proxy service (**Bypass/Intercept**), you can do so from the main **Proxy Services** tab. You do not need to enter New/Edit mode to change this setting.


Before you begin, you must understand the goal of your deployment, how the application proxy operates, and the IP addresses (source and/or destination) and ports to intercept. Some proxy services, such as **DNS**, are simple—comprised only of IP addresses and ports. Others, such as **HTTP**, have more attributes to consider.

For a high-level description of these options, see "[About Proxy Attributes in the Services](#)" on page 114.


For specific proxy descriptions, see

To create a new proxy service:

1. From the Management Console, select the **Configuration > Services > Proxy Services** tab.



2. At the bottom of the tab, click **New Service Group**. The **New Service Group** dialog displays.
3. In the **Service Group** field, name the custom service.
4. Click **OK**. The new service displays under **Custom Service Groups**.



5. Click **New Service**. The New Service dialog displays.

6. Configure service attributes, including applicable proxy settings:
 - a. In the **Name** field, enter a name that describes the service.
 - b. From the **Service Group** drop-down list, select which group displays the service on the main page. You can add the service to a default group or any already-created custom groups.
 - c. **Proxy Settings**—From the **Proxy** drop-down list, select the supported proxy that is compatible with the application protocol.

The **Proxy Settings** sub-options are dynamic (including **TCP/IP Settings**), based on the selected proxy. See ["About Proxy Attributes in the Services"](#) on page 114 for overviews of these options; for more detailed information, see the chapter that explains each proxy in more detail.

Note: The **Detect Protocol** setting is disabled by default. You must select this check box for filtering to be recognized.

- d. **Application Delivery Network Settings**—(Not available for all proxies).

Enable ADN—This setting does not guarantee acceleration for this service—it also depends on ADN routing (for explicit deployments) or network setup (for transparent deployments).

Enable byte caching—This acceleration technique replaces byte sequences in traffic flows with reference tokens and stores them in a byte cache on a pair of ProxySG appliances at each end of the WAN. When a matching byte sequence is requested again, the ProxySG transmits a token instead of the byte sequence.

Enable compression—Uses a variety of algorithms to remove extraneous/predictable information from the traffic before it is transmitted. The information is reconstituted at the destination based on the same algorithms.

Note: To get the maximum benefit of ADN, both byte caching and compression should be enabled. In cases where byte caching may be causing issues for an ADN deployment, you can turn off the **Enable byte caching** option and just use compression (or vice versa). If you know the traffic for this proxy is already compressed or encrypted, you can conserve resources by clearing the **Enable byte caching** and **Enable compression** options. For additional information about byte caching and compression, see "[ADN Acceleration Techniques](#)" on page 713.


Enable thin client processing—Applies special treatment to application traffic from thin client applications (such as RDP, VNC, and Citrix). This processing improves responsiveness of thin client actions. For example, end-users will notice that the desktop displays significantly faster. In addition, thin client data is not retained in the byte cache as long as other types of data because this data is more temporal in nature; the byte cache, therefore, can be used more efficiently for other types of traffic that can better leverage it.

This option is available for TCP Tunnel proxies only, and is only available when ADN is enabled and byte caching and/or compression is enabled. Retention priority and thin client processing are mutually exclusive settings; you cannot enable both options for a service.


Note: For thin client processing to be most effective, you must deactivate the thin client's software-based encryption and compression.

Retention priority—You can control how long data is stored in the byte cache dictionary by assigning a *retention priority* to a particular service. If you want to keep certain types of data in the dictionary for as long as possible, set a **high** retention priority for the service. Or for data that isn't likely to get much benefit from byte caching, you can set a **low** retention priority for the related service. Most services are set to **normal** priority by default. This option is available only if byte caching is enabled for the service.

You can use this option to preserve the most relevant content in the byte cache in the face of continually incoming, competing byte cache data. For example, when an application is being used for backup, you may want to set the retention priority to high so that competing traffic doesn't evict the backup data. However, if an application is being used for data replication, you may want to set the service's retention priority to low as the data most likely will only be hit in the next short duration.



7. Create a listener, or the IP address(es) and ports that this application protocol uses. In the **Listeners** area, click **New**. The **New Listener** dialog displays.



8. Configure the new listener attributes:
 - a. In the **Source address** area, the most common selection is **All**, which means the service applies to requests from any client (IPv4 and IPv6). You can also restrict this listener to a specific IP address (IPv4 or IPv6) or user subnet (for IPv4) or prefix length (for IPv6).
 - b. Select a **Destination address** from the options. The correct selection might depend on network configuration. For overviews of the options, see ["About Proxy Services"](#) on page 110.

- c. In the **Port Range** field, enter a single port number or a port range on which this application protocol broadcasts. For a port ranges, enter a dash between the start and end ports. For example: 8080-8085
 - d. In the **Action** area, select the default action for the service: **Bypass** configures the service to ignore any traffic matching this listener. **Intercept** configures the service to intercept and proxy the associated traffic.
 - e. Click **OK** to close the dialog. The new listener displays in the **Listeners** area.
- 9. Click **Ok** add the new service to the selected service group.
 - 10. Click **Apply**.

See Also

- ❑ ["Moving a Service"](#)
- ❑ ["Importing a Service from the Service Library"](#)

Section D: Proxy Service Maintenance Tasks

This section provides various tasks for managing existing services.

- ["Moving a Service"](#)
- ["Deleting a Service or Service Group" on page 127](#)
- ["Bypassing All Proxy Services \(Troubleshooting\)" on page 128](#)
- ["Importing a Service from the Service Library" on page 128](#)


Moving a Service

The predefined services are not anchored to their default groups. You can move a service to any other predefined or custom group.

Note: You must move the entire service; that is, you cannot move individual service listeners.

To move a service to another service group:

1. From the Management Console, select the **Configuration > Services > Proxy Services** tab.



2. Move the service:
 - a. Select a service.
 - b. Click **Move Service**. The Move Service dialog displays.
 - c. From the drop-down list, select an existing service group (custom or pre-defined).
 - d. Click **OK**.
3. Click **Apply**.

Deleting a Service or Service Group

You can delete a service within a predefined service group but you cannot delete an empty predefined service group itself. However, you can delete a custom service group if it is empty.

You can add back a default service you deleted from the service library by using the Import Service feature. See ["Importing a Service from the Service Library"](#) on page 128.

To delete a service:

1. From the Management Console, select the **Configuration > Services > Proxy Services** tab.
2. Select the service or custom service group to delete.
3. Click **Delete**. A confirmation prompt displays.
4. Click **Yes**. The selected service or custom service group is deleted.
5. Click **Apply**.


Bypassing All Proxy Services (Troubleshooting)

The Bypass All Proxies feature is intended as an interim solution while application-breaking problems are repaired. When **Force Bypass** is invoked, transparent proxy connections are bypassed and explicit proxy connections are rejected.

Note: Downgrading to a version that does not support force bypass while running in bypass mode will result in restoration of proxy services.

To bypass all proxy services:

1. From the Management Console, select the **Configuration > Services > Proxy Services** tab.



2. In the **Force Bypass** area, select the **Temporarily bypass all proxy services** option. The bypass statement to red.
3. Click **Apply**.


Importing a Service from the Service Library

Importing a service procedure is required if you delete a default service and want to re-add it. If you import an existing service, you are prompted to confirm the replacement of a service. Existing service settings are overwritten with the default settings.

In addition, after upgrading the software, any new services added to the service library must be imported if you want to use them.

To import a service from the service library:

1. From the Management Console, select the **Configuration > Services > Proxy Services > Proxy Services** tab.
2. Click **Import Service**. The **Import Service** dialog displays.



3. Configure the import service options:
 - a. From the **Name** drop-down list, select the service to import.
 - b. All other settings adjust automatically to the service's default values. Perform changes if required.
 - c. Click **New** to configure a new listener or **Edit** to modify existing listener settings.
 - d. Click **OK**.
4. Click **Apply**.

Section E: Global Options for Proxy Services

This section describes features that apply to all proxies and services. See ["Proxy Service Global Options"](#) for details.

Proxy Service Global Options

Blue Coat provides certain option settings that when configured apply to *all* proxy services:

- ❑ ["Ensuring Application Availability \(Tunnel on Protocol Error\)"](#)
- ❑ ["Using the Client IP Address for Server Connections"](#) on page 132
- ❑ ["Improving Performance by Not Performing a DNS Lookup"](#) on page 133
- ❑ ["Managing Licensed User Connection Limits \(ProxySG to Server\)"](#) on page 136

Note: You can subscribe to the CachePulse service to optimize HTTP traffic. For information, see ["Enabling CachePulse"](#) on page 182.

Ensuring Application Availability (Tunnel on Protocol Error)

HTTP Proxy

In many networks, business-critical applications send traffic over port 80—the default HTTP port—because it is used as a generic route through the firewall. However, the ProxySG HTTP proxy encounters problems when it receives non-HTTP requests from clients or browsers. The client receives an exception page and the connection closes. The following deployment operations create this situation:

- ❑ The client request from an application or browser is not HTTP.
- ❑ The request is HTTP but also contains components that are not HTTP.
- ❑ The request contains an unexpected formatting error in a line or header.

The ProxySG provides an option that enables the HTTP proxy to tunnel the connection when it receives non-HTTP traffic or broken HTTP request. This allows application traffic to continue and employee production to continue. The transactions remain labeled as HTTP; therefore, the access logs and the **Traffic Mix** and **Active Sessions** statistics display **TCP_TUNNELED** to indicate when a connection passed through the HTTP proxy. The HTTP proxy cannot apply security policies; however, benefits provided by ADN configurations might occur.

The TCP Tunnel on Error option is viable with the following deployments:

- ❑ Applies only to HTTP traffic; HTTPS is not supported in either forward or reverse proxy modes.
- ❑ Applies only to errors in requests from the client browser or application to the ProxySG. Any issues that arise from server responses are not accommodated by this feature.

SSL Proxy

For the SSL proxy, the Tunnel on Protocol Error option applies when non-SSL traffic arrives at the SSL port (443 by default). A common scenario that causes this is having peer-to-peer applications (viz, Skype, BitTorrent, Gnutella, older AOL-IM and eMule) configured to enable port 443 for peer-to-peer traffic without SSL set as the transport protocol. A ProxySG transparently intercepting all 443 traffic cannot process these connections, rendering the application unusable.

With an explicit proxy deployment, SSL errors during the initial handshake causes the same issue. The following example illustrates this:

- ❑ ProxySG is configured to have an explicit HTTP service on port 8080.
- ❑ The HTTP service is configured with *detect protocol* enabled, which hands off SSL traffic to the SSL proxy from an `HTTP CONNECT` request. Detect Protocol is set to OFF by default.

Note: The same applies to an explicit SOCKS proxy deployment with protocol detection enabled or an explicit TCP listener.


Forwarding Note

Enabling the TCP Tunnel on Error option might cause issues if the ProxySG has forwarding rules that direct traffic to upstream proxies or other devices:

- ❑ Forwarding hosts are not viewed as HTTP proxies (even if they are). The initial ProxySG HTTP proxy connects with a TCP tunnel to the forwarding host. If the ProxySG has a policy to forward and tunnels on error, the forwarding rule might not match if the forwarding rule has a condition based on information that is not present—any HTTP conditions, such as:
 - Request method
 - Request URL
 - Request headers
- ❑ In the case of tunnel on error with explicit proxy, HTTP must match a forwarding host for the connection of a successful TCP tunnel to occur. If no forwarding host matches, HTTP will not tunnel on error.

To enable TCP tunnel on HTTP protocol errors:

1. Select the **Configuration > Proxy Settings > General > General** tab.



2. In the **Tunnel on Protocol Error** area, select **TCP tunnel requests when a protocol error is detected**.
3. Click **Apply**.

Related Policy

The Visual Policy Manager (VPM) provides the **Client Certificate Requested** object in the **SSL Intercept Layer > Service** column (the equivalent CPL is `client.certificate.requested={yes|no}`). Use this policy in conjunction with an `SSL.Intercept(no)` action, or a **Do Not Intercept SSL** action in the VPM, to minimize traffic disruption when the SSL proxy intercepts secure traffic where the OCS requests a client certificate.

When Tunnel on Error is enabled, the first detection of a client certificate request from an OCS causes the connection to fail. The appliance adds the details for that exchange to an internal list of connections for which SSL interception should be negated. Subsequent requests function as expected.

Using the Client IP Address for Server Connections

This section discusses configuring the ProxySG to use the IP address of the client to connect to destination servers rather than use the ProxySG address.

About Reflecting the Client Source IP when Connecting to Servers

By default, the ProxySG uses its own IP address as the source IP address for requests (when connecting to servers). If Reflect Client IP is enabled, the ProxySG uses the *client* IP address for all requests. Enabling this option is not an arbitrary decision; it depends on the deployment and role of the ProxySG. For example, if this ProxySG is acting as a branch peer in an Application Delivery Network (ADN) deployment, enable client IP address reflection. This provides maximum visibility for network usage statistics and enables user-based access control to network resources.


Note: The **Reflect Client IP** option is only supported in transparent ProxySG deployments.

You can globally enable the Reflect Client IP option for all services that will be intercepted. To apply Reflect Client IP option to only a few services, first enable this option globally and then create policy to disable the Reflect Client IP option for the exceptions. Or, disable the option globally and create policy to enable it.

Enabling Reflect Client Source IP

To configure the ProxySG to connect to servers using client source IP addresses:

1. Select the **Configuration > Proxy Settings > General > General** tab.



2. In the **Reflect Client IP** area, select **Reflect client's source IP when connecting to servers**.
3. Click **Apply**.

Important: If you enable Reflect Client IP and want the ProxySG to preserve persistent client connections, you must also add policy.

VPM object: **Web Access Layer > Action > Support Persistent Client Requests** (static)

CPL:

```
<proxy>  
    http.client.persistence(preserve)
```

Improving Performance by Not Performing a DNS Lookup

This section describes how to improve performance by configuring the ProxySG to trust the destination IP address provided by the client.

About Trusting the Destination IP Address Provided by the Client

If, in your environment, a client sometimes provides a destination IP address that the ProxySG cannot identify, you have the option to configure the ProxySG to *not* perform a DNS lookup and allow that IP address. This can improve performance, but potentially presents a security issue.

You can configure the ProxySG appliance to trust a client-provided destination IP address in transparent proxy deployments where:

- ❑ DNS configuration on the client is correct, but is not correct on the ProxySG.

- ❑ The client obtains the destination IP address using Windows Internet Name Service (WINS) for NetBIOS name resolution.
- ❑ DNS imputing on the ProxySG is not configured correctly. On the ProxySG, you can configure a list of suffixes to help with DNS resolution. In the event that the host name is not found, these suffixes are appended to the host name provided by the client. For information on DNS imputing, see ["Resolving Hostnames Using Name Imputing Suffixes"](#) on page 819.

In each of the cases above, the ProxySG cannot obtain the destination IP address to serve client requests. When you enable the ProxySG to trust a client-provided destination IP address, the ProxySG uses the IP address provided by the client and does not perform a DNS lookup.




Figure 7–1 No DNS lookup occurs; the transactions goes straight to the OCS.





Figure 7–2 The ProxySG initiates a DNS lookup and initiates a new connection to the server. The ProxySG cannot trust the client-provided destination IP address in the following situations:

- ❑ The ProxySG receives the client requests in an explicit proxy deployment.
- ❑ The ProxySG has a forwarding rule configured for the request.
- ❑ The ProxySG has a SOCKS gateway rule configured for the request.
- ❑ The ProxySG has policy that rewrites the server URL.

A transproxy deployment is one where a client is configured to contact a ProxySG explicitly, and a new ProxySG is deployed between the client and its explicit proxy. The new ProxySG, now transparently intercepts the traffic between the client and its explicit proxy. In a transproxy deployment, the destination IP address used by the client does not match the host header in the HTTP request, since the client is configured to use the explicit proxy. The path that the client request takes in a transproxy deployment depends on whether or not **Trust Destination IP** is enabled on the transparently deployed ProxySG.

- ❑ When **Trust Destination IP** is enabled on the transparent ProxySG, the transparent proxy trusts the destination IP included in the request and forwards the request to the explicit proxy which is serviced either from cache or from the Origin Content Server (OCS).
- ❑ When **Trust Destination IP** is disabled on the transparent ProxySG, the transparent proxy performs a DNS resolution on the host header in the request. The request is then completed based on the configured policy—forwarding rules, SOCKS gateway policy, and server URL rewrite policy.



Note: If a client gives the destination address of a blocked site but the host name of a non-blocked site, with Trust Destination IP enabled, the ProxySG connects to the destination address. This might allow clients to bypass the configured security policy for your environment.

About the Default Settings

During the ProxySG initial configuration tasks, the administrator determined the default Trust Destination IP setting. In most deployments, the role of the ProxySG determines the setting:

- ❑ Acceleration role: enabled.
- ❑ Most other proxy deployments: disabled for tighter security.


You can change these defaults through the Management Console, the CLI, or through policy. If you use policy, however, be aware that it overrides the setting in the in Management Console.

For information about using the `trust_destination_ip(yes|no)` CPL property, refer to the *Content Policy Language Guide*.

Configuring the ProxySG to Trust or Not Trust the Destination IP Address

To change the current trust destination default setting:

1. Select the **Configuration > Proxy Settings > General** tab.



2. Select or clear the **Trust client-provided destination IP when connecting to servers** option.
3. Click **Apply**.

Managing Licensed User Connection Limits (ProxySG to Server)

This section describes ProxySG how to enable license-enforced user limits, describes how to monitor user numbers, and describes how to configure the ProxySG to behave when a limit is breached.

About User Limits

If you have more users connecting through the system than is coded by the model license, you have an option to configure the overflow behavior (after a permanent model license has been applied to the system). The enforcement options are queue the connections or bypass through ProxySG and proceed directly to the server.

Only unique IP addresses of connections intercepted by proxy services are counted toward the user limit; furthermore, the number of users depends on the hardware model and whether or not ADN is enabled.

License-enforced user connection limits are *advisory* and are based on optimal performance for each ProxySG. The default setting is to *not* enforce user limits; however, when a user connection limit is breached, the ProxySG logs the event and the license health indicator changes to **Critical**.

For WAN optimization deployments, Symantec recommends purchasing a ProxySG model based on the maximum number of client connections it needs to support, not the maximum number of users, since the connection limit is likely to be reached first; your channel partner SE or local Symantec SE can assist you with WAN optimization connection counts and sizing for your specific needs.

The following tables provides the user connection limits hard-coded into the license per ProxySG and ProxySG VA model.

Table 7–4 Hardware Models and Licensed Users

ProxySG Model	Number of Licensed Users (Concurrent Source IP Addresses)	
	Without ADN	With ADN Enabled
210-5	30	10
210-10	150	50
210-25	Not License Limited	Not License Limited
300-5	30	10
300-10	150	150
300-25	Not License Limited	Not License Limited
510-5	200	50
510-10	500	125
510-20	1200	300
510-25	Not License Limited	Not License Limited
600-10	500	500
600-20	1000	1000
600-35	Not License Limited	Not License Limited
810-5	2500	500
810-10	3500	700
810-20	5000	1000
810-25	Not License Limited	Not License Limited
900-10	3500	3500
900-20	6000	6000
900-30, 900-45	Not License Limited	Not License Limited
8100-5, 8100-10, 8100-20, 8100-20-DC	Not License Limited	Not License Limited

Table 7–4 Hardware Models and Licensed Users (Continued)

ProxySG Model	Number of Licensed Users (Concurrent Source IP Addresses)	
9000-5, 9000-10, 9000-20, 9000-30, 9000-40	Not License Limited	Not License Limited

Table 7–5 Virtual Appliance Models and Licensed Users

ProxySG VA Model	Number of Licensed Users
VA-5	10
VA-10	50
VA-15	125
VA-20	300

Tasks for Managing User Limits

To learn more about user limits, see ["About User Limits"](#) on page 136.

Monitoring and managing user limits requires the following tasks:

- ❑ ["Modifying User Limits Notifications"](#) on page 138—Configure the ProxySG to monitor and alert you when a user limit is near.
- ❑ ["Determining Behavior When User Limits are Exceeded"](#) on page 140—Determine what happens when more user connections than allowed by the license occurs.


Modifying User Limits Notifications

You can set and monitor user limit thresholds of the model license. A threshold breach triggers a notification and/or event log entry. Frequent breaches indicate that constant user connections to this particular ProxySG model are exceeding the optimal design.


Note: You can access the **Statistics > Health Monitoring > Licensing** tab to view licensing status, but you cannot make changes to the threshold values from that tab.

To view licensing metrics and set user limits notifications:

1. Click **Maintenance > Health Monitoring > Licensing**.



2. Select **User License Utilization**.
3. Click **Edit**. The Edit Health Monitor Settings dialog displays.



4. (Optional) Modify the threshold and interval values to your satisfaction. The thresholds represent the percentage of license use.
 - a. Modify the **Critical** and/or **Warning Threshold** settings. These values are the percentages of maximums. For example, if the ProxySG is an SG810-20 and ADN is enabled, the maximum number of unique users connections is 1000. With a **Warning Threshold** value of **80** (percent) and **Critical Threshold** value of **90**, the notification triggers when user connectivity reaches **800** and **900**, respectively.
 - b. Modify the **Critical** and/or **Warning Interval** settings. These values are the number of seconds that elapse between user limit checks. By default, both critical and warning interval checks occur every **120** seconds.
5. Select the notification settings:
 - **Log** adds an entry to the Event Log.
 - **Trap** sends an SNMP trap to all configured management stations.
 - **Email** sends an e-mail to the addresses listed in the Event Logging properties (**Maintenance > Event Logging > Mail**).
6. Click **OK** to close the dialog.
7. Click **Apply**.

For information about licensing, see "[Licensing](#)" on page 43.


Determining Behavior When User Limits are Exceeded

You can specify what happens when more users simultaneously connect through the ProxySG (*overflow connections*) than is allowed by the model license:

- ❑ Bypass the system: All connections exceeding the maximum are passed through the system without processing.
- ❑ Queue connections: All connections exceeding the maximum are queued, waiting for another connection to drop off.
- ❑ Do not enforce the licensed user limit: This is the default option for hardware appliances. This allows for unlimited connections; however, exceeding the license limit triggers a health monitoring event. This option is not available for virtual appliances because the ProxySG VA always enforces the licensed user limit.

To specify what happens when overflow connections occur:


1. Select **Configuration > Proxy Settings > General**.



2. In the **User Overflow Action** area, select an action that occurs when the licensed user limits are exceeded:
 - **Do not enforce licensed user limit** is the default. Unlimited user connections are possible. If the limit is exceeded, the ProxySG health changes to **CRITICAL**. This option is not available on the ProxySG VA because licensed user limits are always enforced.
 - **Bypass connections from users over license limit**—Any transaction from a user whose connection exceeds the licensed limit is not susceptible to policy checks or any other ProxySG benefit, such as acceleration. This option provides the best user experience (with the caveat of potentially slower performance), but presents a Web security risk. This is the default option for the ProxySG VA.
 - **Queue connections from users over license limit**—Any transaction from a user whose connection exceeds the licensed limit must wait (in order) for an available ProxySG connection. This option provides the lowest user experience (and users might become frustrated and, perceiving a hang, might attempt request refreshes), but preserves Web security policies.
3. Click **Apply**.

Viewing Concurrent Users

View a snapshot of intercepted, concurrent users by selecting the **Statistics > System > Resources > Concurrent Users** tab. The tab shows user connections going through the ProxySG for the last 60 minutes, day, week, month, and year. Only unique IP addresses of connections intercepted by proxy services are counted toward the user limit.



See Also

- ❑ "Global Options for Proxy Services"
- ❑ "Enabling Reflect Client Source IP"
- ❑ "About Trusting the Destination IP Address Provided by the Client"
- ❑ "Managing Licensed User Connection Limits (ProxySG to Server)"

Section F: Exempting Requests From Specific Clients

The bypass list contains IP addresses/subnet masks of client and server workstations. Used only in a transparent proxy environment, the bypass list allows the ProxySG to skip processing requests sent from specific clients to specific servers. The list allows traffic between protocol incompliant clients and servers to pass through the ProxySG without a disruption in service.

Note: This prevents the appliance from enforcing any policy on these requests and disables any caching of the corresponding responses. Because bypass entries bypass Blue Coat policy, use bypass sparingly and only for specific situations.

This section covers the following topics:

- ❑ ["Adding Static Bypass Entries"](#)
- ❑ ["Using Policy to Configure Dynamic Bypass"](#) on page 144


Adding Static Bypass Entries

You can add entries to prevent the ProxySG from intercepting requests from specified systems.

Note: Dynamic bypass cannot be configured through the Management Console. You must define policy or use the CLI. For more information, see ["Using Policy to Configure Dynamic Bypass"](#) on page 144.

To add static bypass entries:

1. Click the **Configuration > Services > Proxy Services > Static Bypass List** tab.



2. Click **New** to create a new list entry (or click **Edit** to modify a list entry). The New Bypass List Entry dialog displays.
3. Create a **Client Address** or **Server Address** entry. The IP address can be IPv4 or IPv6. If you enter an IPv4 address, you can specify a subnet mask. For IPv6 addresses, you can specify a prefix length.
4. (Optional) Add a Comment that indicates why you are creating the static bypass rule for the specific source/destination combination. This is useful if another administrator needs to tune the settings later.
5. Click **OK** to close the dialog.
6. Click **Apply**.

Using Policy to Configure Dynamic Bypass

Dynamic bypass, available through policy, can automatically compile a list of response URLs that return various types of errors.

Note: Because bypass entries bypass Blue Coat policy, the feature should be used sparingly and only for specific situations.

About Dynamic Bypass

Dynamic bypass keeps its own (dynamic) list of which connections to bypass, where connections are identified by both source and destination. Dynamic bypass can be based on any combination of policy triggers. In addition, some global settings can be used to selectively enable dynamic bypass based on specific HTTP response codes. After an entry exists in the dynamic bypass table for a specific

source/destination IP pair, all connections from that source IP to that destination IP are bypassed in the same way as connections that match against the static bypass list.

For a configured period of time, further requests for the error-causing URLs are sent immediately to the origin content server (OCS), bypassing the ProxySG. The amount of time a dynamic bypass entry stays in the list and the types of errors that cause the ProxySG to add a site to the list, as well as several other settings, are configurable from the CLI.

After the dynamic bypass timeout for a client and server IP address entry ends, the ProxySG removes the entry from the bypass list. On the next client request for the client and server IP address, the ProxySG attempts to contact the OCS. If the OCS still returns an error, the entry is again added to the local bypass list for the configured dynamic bypass timeout. If the entry does not return an error, entries are again added to the dynamic list and not the local list.

Notes

- ❑ Dynamic bypass entries are lost when the ProxySG is restarted.
- ❑ No policy enforcement occurs on client requests that match entries in the dynamic or static bypass list.
- ❑ If a site that requires forwarding policy to reach its destination is entered into the bypass list, the site is inaccessible.

Configuring Dynamic Bypass

Dynamic bypass is disabled by default. Enabling and fine-tuning dynamic bypass is a two-step process:

- ❑ Set the desired dynamic bypass timeout and threshold parameters.
- ❑ Use policy (recommended) or the CLI to enable dynamic bypass and set the types of errors that cause dynamic bypass to add an entry to the bypass list.

Adding Dynamic Bypass Parameters to the Local Bypass List

The first step in configuring dynamic bypass is to set the `server-threshold`, `max-entries`, or `timeout` values in the CLI.

Note: This step is optional because the ProxySG uses default configurations if you do not specify them. Use the default values unless you have specific reasons for changing them. Contact Blue Coat Technical Support for detailed advice on customizing these settings.

- ❑ The `server-threshold` value defines the maximum number of client entries before the ProxySG consolidates client-server pair entries into a single server entry that then applies to all clients connecting to that server. The range is 1 to 256. The default is 16. When a consolidation occurs, the lifetime of the consolidated entry is set to the value of `timeout`.

- ❑ The `max-entries` defines the maximum number of total dynamic bypass entries. The range is 100 to 50,000. The default value is 10,000. When the number of entries exceeds the `max-entries` value, the oldest entry is replaced by the newest entry.
- ❑ The `timeout` value defines the number of minutes a dynamic bypass entry can remain un referenced before it is deleted from the bypass list. The range is 1 to 86400. The default value is 60.

Enabling Dynamic Bypass and Specifying Triggers

Enabling dynamic bypass and specifying the types of errors that causes a URL to be added to the local bypass list are done with the CLI. You cannot use the Management Console.

Using policy to enable dynamic bypass and specify trigger events is better than using the CLI, because the CLI has only a limited set of responses. For information about available CLI triggers, refer to the *Blue Coat Content Policy Language Guide*. For information about using policy to configure dynamic bypass, refer to the *Visual Policy Manager Reference*.

Bypassing Connection and Receiving Errors

In addition to setting HTTP code triggers, you can enable connection and receive errors for dynamic bypass.

If `connect-error` is enabled, any connection failure to the origin content server (OCS), including timeouts, inserts the OCS destination IP address into the dynamic bypass list.

If `receive-error` is enabled, when the cache does not receive an HTTP response on a successful TCP connection to the OCS, the OCS destination IP address is inserted into the dynamic bypass list. Server timeouts can also trigger `receive-error`. The default timeout value is 180 seconds, which can be changed.

CLI Syntax to Enable Dynamic Bypass and Trigger Events

- ❑ To enter configuration mode for the service:

```
SGOS#(config) proxy-services
SGOS#(config proxy-services) dynamic-bypass
```

- ❑ The following subcommands are available:

```
SGOS#(config dynamic-bypass) {enable | disable}
SGOS#(config dynamic-bypass) max-entries number
SGOS#(config dynamic-bypass) server-threshold number
SGOS#(config dynamic-bypass) trigger {all | connect-error | non-http |
receive-error | 400 | 401 | 403 | 405 | 406 | 500 | 502 | 503 | 504}
SGOS#(config dynamic-bypass) timeout minutes
SGOS#(config dynamic-bypass) no trigger {all | connect-error | non
http | receive-error | 400 | 401 | 403 | 405 | 406 | 500 | 502 | 503 |
504}
SGOS#(config dynamic-bypass) clear
SGOS#(config dynamic-bypass) view
```


Section G: Trial or Troubleshooting: Restricting Interception From Clients or To Servers

This section discusses Restricted Intercept topics. See ["Restricted Intercept Topics"](#) for details.

Restricted Intercept Topics

- ❑ ["About Restricted Intercept Lists"](#)
- ❑ ["Creating a Restricted Intercept List"](#) on page 147

About Restricted Intercept Lists

By default, all clients and servers evaluate the entries in Proxy Services where the decision is made to intercept or bypass a connection. To restrict or reduce the clients and servers that can be intercepted by proxy services, create *restricted intercept lists*. A restricted intercept list is useful in a rollout, before entering full production—you only want to intercept a subset of the clients. After the ProxySG is in full production mode, you can disable the restricted intercept list.

A restricted intercept list is also useful when troubleshooting an issue because you can reduce the set of systems that are intercepted.


Notes

- ❑ Restricted intercepts lists are only applicable to transparent connections.
- ❑ An entry can exist in both the **Static Bypass List** and the **Restricted Intercept List**. However, the **Static Bypass List** overrides the entries in the **Restricted Intercept List**.

Creating a Restricted Intercept List

To create a Restricted Intercept List:

1. From the Management Console, select the **Configuration > Services > Proxy Services > Restricted Intercept List** tab.



2. Select **Restrict Interception to the servers and clients listed below-- all other connections are bypassed**.
3. Create a new entry:
 - a. Click **New**; the New Restricted Intercept Entry dialog displays.
 - b. Restrict interception from specific clients: In the **Client Address** area, select **Client host or subnet**. Enter an IPv4 or IPv6 address in the **IP Address** field and enter the subnet mask (for IPv4 addresses) or prefix length (IPv6) in the **Prefix/Subnet** field.
 - c. Restrict interception to specific servers: In the **Server Address** area, select **Server host or subnet**. Enter an IPv4 or IPv6 address in the **IP Address** field and enter the subnet mask (for IPv4 addresses) or prefix length (IPv6) in the **Prefix/Subnet** field.
 - d. Click **OK** to close the dialog.
4. Click **Apply**.

Section H: Reference: Proxy Services, Proxy Configurations, and Policy

This section provides reference material.

- ❑ ["Reference: Proxy Types"](#)
- ❑ ["Reference: Service/Proxy Matrices"](#) on page 151
- ❑ ["Reference: Access Log Fields"](#) on page 152

Reference: Proxy Types

This section provides descriptions of the available proxies.

Table 7–6 Proxy Types

Proxy Name	Protocol/Description	Capabilities and Benefits
AOL-IM	AOL Instant Messaging	- Controls AOL instant messaging actions by allowing or denying IM communications and file sharing based on users (both employee identities and IM handles), groups, file types and names, and other triggers. - All IM communications can be logged and archived for review.
CIFS	Common Internet File System	Optimizes/accelerates file sharing across the WAN to users in branch offices.
DNS	Domain Name Service	- Speeds up domain name resolution by looking up domain names in the ProxySG appliance's DNS cache. If the name isn't found in the cache, the ProxySG forwards the request to the configured DNS server list. - Ability to rewrite DNS requests and responses.
Flash	Adobe Flash Real Time Messaging Protocol	- Live streaming—The ProxySG appliance fetches the live Flash stream <i>once</i> from the OCS and serves it to <i>all</i> users behind the appliance. - Video-on-demand—As Flash clients stream pre-recorded content from the OCS through the ProxySG, the content is cached on the appliance. After content gets cached on the ProxySG, subsequent requests for the cached portions are served from the appliance; uncached portions are fetched from the OCS.
FTP	File Transfer Protocol	- Controls, secures, and accelerates file transfer requests - Caches FTP objects.
HTTP	Hyper Text Transfer Protocol	- Controls, secures, and accelerates Web traffic - Caches copies of frequently requested web pages and objects.

Table 7–6 Proxy Types (Continued)

Proxy Name	Protocol/Description	Capabilities and Benefits
HTTPS Reverse Proxy	A proxy positioned in front of an HTTPS server that answers secure web requests from clients (using the ProxySG appliance's local cache when possible)	- Accelerates secure web requests, improving the response time to clients. - Because the Reverse Proxy is processing the requests, it allows the HTTPS server to handle a heavier traffic load.
MAPI	Messaging Application Programming Interface; protocol used by Microsoft Outlook (client) to communicate with Microsoft Exchange (server).	Accelerates the following Outlook processes: sending/receiving e-mail, accessing message folders, changing calendar elements.
MMS	Microsoft Media Services; streaming protocol	- Monitors, controls, limits, or blocks streaming media traffic that uses Microsoft's proprietary streaming protocol. - Reduces stutter and improves the quality of streaming media. - Logs streaming connections.
MSN-IM	MSN Instant Messaging	- Controls MSN instant messaging actions by allowing or denying IM communications and file sharing based on users, groups, file types and names, and other triggers. - Logs all IM communications for review.
RTSP	Real Time Streaming Protocol	- Monitors, controls, limits, or blocks streaming media traffic that uses the Internet standard RTSP protocol. - Reduces stutter and improves the quality of streaming media. - Logs streaming connections.
Shell	A proxy that allows a client to connect to other destinations via Telnet, after the client has created an authenticated Telnet connection to the ProxySG	- Monitors, controls, limits, or blocks outbound Telnet connections. - Enforces access control to a group of users and destinations via policy. - Logs all connections.

Table 7–6 Proxy Types (Continued)

Proxy Name	Protocol/Description	Capabilities and Benefits
SOCKS	A proxy that allows a client to connect to other destination servers/ports in a SOCKS tunnel, after the client's connection to the SOCKS proxy is authenticated	- Monitors, controls, limits, or blocks outbound client connections requested using the SOCKS protocol. - Through policy, enforces access control to a group of users and destinations. - SOCKS traffic can be passed to other proxies (such as HTTP or AOL-IM) for acceleration. - Logs all connections.
SSL	Secure Socket Layer	- Allows authentication, virus scanning and URL filtering of encrypted HTTPS content. - Accelerates performance of HTTPS content, using HTTP caching. - Validates server certificates presented by various secure websites at the gateway.
TCP-Tunnel	A tunnel for any TCP-based protocol for which a more specific proxy is not available	Compresses and accelerates tunneled traffic.
Yahoo-IM	Yahoo Instant Messaging	- Controls Yahoo instant messaging actions by allowing or denying IM communications and file sharing based on users, groups, file types and names, and other triggers. - Logs all unencrypted IM communications for review.

Reference: Service/Proxy Matrices

Expanding on the service port listing at the beginning of this chapter, the table below provides a list of the pre-defined proxy services and listeners that the Proxy can accelerate and interpret. Links to the related proxy configuration sections are included.

Table 7–7 Proxy Name and Listeners (alphabetical order)

Service Name	Proxy	Destination IP Address	Port Range	Configuration Discussed
AOL-IM	AOL-IM	All	5190	
CIFS	CIFS	Transparent	445, 139	"Accelerating File Sharing" on page 255
DNS	DNS	All	53	"Managing the Domain Name Service (DNS) Proxy" on page 301
Endpoint Mapper	Endpoint Mapper	All	135	Chapter 11: "Managing Outlook365 Applications" on page 271

Table 7–7 Proxy Name and Listeners (alphabetical order) (Continued)

Service Name	Proxy	Destination IP Address	Port Range	Configuration Discussed
Explicit HTTP	HTTP	Explicit	8080, 80	"Intercepting and Optimizing HTTP Traffic" on page 153
External HTTP	HTTP	Transparent	80	
FTP	FTP	All	21	"Managing the File Transport Protocol (FTP) Proxy" on page 291
HTTPS	SSL	All	443	"Managing the SSL Proxy" on page 213
Internal HTTP	TCP-Tunnel	192.168.0.0/16 10.0.0.0/8 172.16.0.0/16 169.254.0.0/16 192.0.2.0/24	80	"Intercepting and Optimizing HTTP Traffic" on page 153
MMS	MMS	All	1755	"Managing Streaming Media" on page 529
MSN-IM	MSN-IM	All	1863, 6891	
MS Terminal Services	TCP-Tunnel	Transparent	3389	"Managing Streaming Media" on page 529
SOCKS	SOCKS	Explicit	1080	"Managing a SOCKS Proxy" on page 305
Yahoo-IM	Yahoo-IM	All	5050, 5101	

Reference: Access Log Fields

The access log has two fields: service name and service group name.

- Name of the service used to intercept this connection:

- `x-service-name` (ELFF token) `service.name` (CPL token)

Note: The `x-service-name` field replaces the `s-sitename` field. The `s-sitename` field can still be used for backward compatibility with squid log formats, but it has no CPL equivalent.

- Service group name:

- `x-service-group` (ELFF token) `service.group` (CPL token)

Note: See [Chapter 29: "Creating Custom Access Log Formats"](#) on page 649 and [Chapter 30: "Access Log Formats"](#) on page 657 for detailed information about creating and editing log formats.

Chapter 8: Intercepting and Optimizing HTTP Traffic

This chapter describes how to configure the HTTP proxy to manage traffic and accelerate performance in your environment.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ [Section A: "About the HTTP Proxy" on page 155](#)
- ❑ [Section B: "Changing the External HTTP \(Transparent\) Proxy Service to Intercept All IP Addresses on Port 80" on page 157](#)
- ❑ [Section C: "Managing the HTTP Proxy Performance" on page 158](#)
- ❑ [Section D: "Selecting an HTTP Proxy Acceleration Profile" on page 173](#)
- ❑ [Section E: "Using a Caching Service" on page 182](#)
- ❑ [Section F: "Fine-Tuning Bandwidth Gain" on page 184](#)
- ❑ [Section G: "Caching Authenticated Data \(CAD\) and Caching Proxy Authenticated Data \(CPAD\)" on page 191](#)
- ❑ [Section H: "Viewing HTTP/FTP Statistics" on page 202](#)
- ❑ [Section I: "Supporting IWA Authentication in an Explicit HTTP Proxy" on page 207](#)
- ❑ [Section J: "Supporting Authentication on an Upstream Explicit Proxy" on page 209](#)
- ❑ [Section K: "Detect and Handle WebSocket Traffic" on page 210](#)

How Do I...?

To navigate this chapter, identify the task to perform and click the link:

How do I...?	See...
Intercept traffic on the HTTP Proxy?	"Changing the External HTTP (Transparent) Proxy Service to Intercept All IP Addresses on Port 80" on page 157
Create a new HTTP Proxy service?	Section C: "Creating Custom Proxy Services" on page 120
Configure the HTTP Proxy for object freshness?	"Allocating Bandwidth to Refresh Objects in Cache" on page 184 Step 4 in "To set HTTP default object caching policy:" on page 171
Bypass the cache or not cache content using policy?	Refer to: <i>Visual Policy Manager Reference</i> <i>Content Policy Language Guide</i> Use either the VPM or CPL to create policy that allows for bypassing the cache or for prohibiting caching based on your needs.
Choose a proxy acceleration profile?	"Selecting an HTTP Proxy Acceleration Profile" on page 173
Cache content without having to use policy?	"Using a Caching Service" on page 182
Configure the HTTP proxy to be a: server accelerator or reverse proxy? forward proxy? server-side bandwidth accelerator?	"About the Normal Profile" on page 173 "About the Portal Profile" on page 173 "About the Bandwidth Gain Profile" on page 174
Fine-tune the HTTP Proxy for bandwidth gain?	"Using a Caching Service" on page 182 "Using Byte-Range Support" on page 185
Configure Internet Explorer to explicitly proxy HTTP traffic?	"Supporting IWA Authentication in an Explicit HTTP Proxy" on page 207
Configure the ProxySG appliance to detect and handle WebSocket traffic?	"Detect and Handle WebSocket Traffic" on page 210

Section A: About the HTTP Proxy

Before Reading Further

Before reading this section, Symantec recommends that you be familiar with the concepts in these sections:

- ❑ "About Proxy Services" on page 110.
- ❑ [Chapter 32: "Configuring an Application Delivery Network"](#) on page 711 (optimize ADN performance on the HTTP Proxy).

The HTTP proxy is designed to manage Web traffic across the WAN or from the Internet, providing:

- ❑ Security
- ❑ Authentication
- ❑ Virus Scanning and Patience Pages
- ❑ Performance, achieved through Object Caching and Object Pipelining
- ❑ Transition functionality between IPv4-only and IPv6-only networks

The proxy can serve requests without contacting the Origin Content Server (OCS) by retrieving content saved from a previous request made by the same client or another client. This is called *caching*. The HTTP proxy caches copies of frequently requested resources on its local hard disk. This significantly reduces upstream bandwidth usage and cost and significantly increases performance.

Proxy services define the ports and addresses where a ProxySG listens for incoming requests. The ProxySG has three default HTTP proxy services: **External HTTP**, **Explicit HTTP**, and **Internal HTTP**. **Explicit HTTP** and **External HTTP** use the HTTP proxy, while **Internal HTTP** uses TCP tunnel.

- ❑ The **Explicit HTTP** proxy service listens on ports 80 and 8080 for explicit connections.
- ❑ The **Internal HTTP** proxy service listens on port 80 and transparently intercepts HTTP traffic from clients to internal network hosts.
- ❑ The **External HTTP** proxy service listens on port 80 for all other transparent connections to the ProxySG. Typically, these requests are for access to Internet resources.

Although you can intercept SSL traffic on either port, to enable the ProxySG to detect the presence of SSL traffic you must enable **Detect Protocol** on the explicit HTTP service so that the SSL traffic is handed off to the SSL Proxy. Default is set to OFF. For more information on SSL proxy functionality, see [Chapter 9: "Managing the SSL Proxy"](#).

Furthermore, you can create a bypass list on the ProxySG to exclude the interception of requests sent from specific clients to specific servers and disable caching of the corresponding responses. The static bypass list also turns off all policy control and acceleration for each matching request. For example, for all clients visiting `www.bluecoat.com` you might exclude interception and caching of

all requests, the corresponding responses, acceleration and policy control. To create a static bypass list, used only in a transparent proxy environment, see ["Adding Static Bypass Entries"](#) on page 143.

When accessing internal IP addresses, Blue Coat recommends using the TCP tunnel proxy instead of the HTTP proxy. Some applications deployed within enterprise networks are not always fully compatible with HTTP specs or are poorly designed. Use of these applications can cause connection disruptions when using HTTP proxy. As a result internal sites and servers use the **Internal HTTP** service, which employs the TCP tunnel proxy.

Important: The TCP tunnel does *not* support HTTP proxy service functionality. That is, only the TCP header of a request, (containing source and destination port and IP) will be visible to the ProxySG for policy evaluation. To ensure you get the most from the appliance, you must edit the External (transparent) HTTP service to use the HTTP proxy instead of the default TCP tunnel.

IPv6 Support

The HTTP proxy is able to communicate using either IPv4 or IPv6, either explicitly or transparently.

In addition, for any service that uses the HTTP proxy, you can create listeners that bypass or intercept connections for IPv6 sources or destinations.

About Web FTP

Web FTP is used when a client uses the HTTP protocol to access an FTP server. Web FTP allows you to connect to a FTP server with the `ftp://` URL. The ProxySG translates the HTTP request into an FTP request for the origin content server (OCS), if the content is not already cached. Further, it translates the FTP response with the file contents into an HTTP response for the client.

To manage Web FTP connection requests on the ProxySG, the HTTP service on port 80 (or 8080 in explicit deployments) must be set to **Intercept**.

For information on using an FTP client to communicate via the FTP protocol, see [Chapter 12: "Managing the File Transport Protocol \(FTP\) Proxy"](#) on page 291.

Configuring Internet Explorer for Web FTP with an Explicit HTTP Proxy

Because a Web FTP client uses HTTP to connect to the ProxySG, the HTTP proxy manages this Web FTP traffic. For an explicitly configured HTTP proxy, Internet Explorer version 6.0 users accessing FTP sites over HTTP must clear the **Enable folder view for FTP sites** browser setting.

To disable Web FTP in Internet Explorer v9.0:

1. In Internet Explorer, select **Tools > Internet Options**.
2. Click the **Advanced** tab.
3. Clear the **Enable FTP folder view** option and click **OK**.


Section B: Changing the External HTTP (Transparent) Proxy Service to Intercept All IP Addresses on Port 80

By default, the **External HTTP** service includes an HTTP proxy service listener configured on port 80. During the initial ProxySG configuration, if it hasn't already been set, you can set **External HTTP** to **Intercept**.

The following procedure describes how to set the service to Intercept mode.

To intercept traffic using the External HTTP proxy service:

1. From the Management Console, select **Configuration > Services > Proxy Services**.



2. Intercept External HTTP traffic:
 - a. Scroll the list of service groups, click **Standard**, and select **External HTTP**.
 - b. Select **Intercept** from the drop-down list.
3. Click **Apply**.

Now that the ProxySG is intercepting HTTP traffic, configure the HTTP proxy options. The following sections provide detailed information and procedures:

- ❑ [Section C: "Managing the HTTP Proxy Performance" on page 158](#)
- ❑ [Section D: "Selecting an HTTP Proxy Acceleration Profile" on page 173](#)
- ❑ [Section E: "Using a Caching Service" on page 182](#)
- ❑ [Section G: "Caching Authenticated Data \(CAD\) and Caching Proxy Authenticated Data \(CPAD\)" on page 191](#)

Section C: Managing the HTTP Proxy Performance

This section describes the methods you can use to configure the HTTP proxy to optimize performance in your network.

- ❑ ["HTTP Optimization"](#)
- ❑ ["Customizing the HTTP Object Caching Policy"](#)
- ❑ ["About the HTTP Object Caching Policy Global Defaults"](#) on page 167
- ❑ ["About Clientless Requests Limits"](#) on page 169
- ❑ ["Preventing Exception Pages From Upstream Connection Errors"](#) on page 170
- ❑ ["Setting the HTTP Default Object Caching Policy"](#) on page 171

HTTP Optimization

The HTTP proxy alleviates the latency in data retrieval and optimizes the delivery of HTTP traffic through object caching and object pipelining. Caching minimizes the transmission of data over the Internet and over the distributed enterprise, thereby improving bandwidth use. Pipelining allows the ProxySG to open several connections to a server, speeding up the delivery of content into the cache. Pre-fetching is another method the ProxySG uses to improve the user experience. Content on a requested web page several levels deep is requested and cached for fast delivery to users.

For objects in cache, an intelligent caching mechanism in the ProxySG maintains object freshness. This is achieved by periodically refreshing the contents of the cache, while maintaining the performance within your network.

The method of storing objects on disk is critical for performance and scalability. SGOS, the operating system on the ProxySG, uses an object store system which hashes object lookups based on the entire URL. This hashing allows access to objects with far fewer lookups, as compared to a directory-based file system found in traditional operating systems. While other file systems run poorly when they are full, the ProxySG's cache system achieves its highest performance when it is full.

Customizing the HTTP Object Caching Policy

Object caching is the saving of an application object locally so that it can be served for future requests without requiring retrieval from the OCS. Objects can, for example, be documents, videos, or images on a Web page. When objects are cached, the only traffic that crosses the WAN are permission checks (when required) and verification checks that ensure that the copy of the object in cache is still fresh. By allowing objects to be shared across requests and users, object caching greatly reduces the bandwidth required to retrieve contents and the latency associated with user requests.

For more information on how the ProxySG executes permission checks to ensure authentication over HTTP, see [Section G: "Caching Authenticated Data \(CAD\) and Caching Proxy Authenticated Data \(CPAD\)"](#) on page 191.

In case of a reverse proxy, object caching reduces the load on the OCS and improves scalability of the OCS.




Figure 8-1 Object Caching on the ProxySG

Before you begin customizing your HTTP Proxy policy, read the following concepts:

- ❑ ["About Object Pipelining" on page 159](#)
- ❑ ["About HTTP Object Freshness" on page 160](#)
- ❑ ["About Meta Tags" on page 161](#)
- ❑ ["About Tolerant HTTP Request Parsing" on page 161](#)
- ❑ ["About HTTP Compression" on page 162](#)
- ❑ ["About the HTTP Object Caching Policy Global Defaults" on page 167](#)

About Object Pipelining

A Web page is typically composed of dozens of objects. When a client requests a Web page, all the objects must be retrieved to display the Web page. This object retrieval process presents a delay for the end user — for example, serial retrieval of the content would create a significant time lag.

Although current browsers open multiple connections with the OCS to retrieve objects in parallel, the ProxySG further accelerates the process with its Object Pipelining algorithm which supports nested pipelines that are up to three levels deep.

- ❑ The Object Pipelining algorithm allows the ProxySG to open as many simultaneous TCP connections as the origin server allows, and retrieves objects in parallel. The proxy also pre-fetches objects based on pipelined requests. For example, if a pipelined HTML object has other embedded objects, the HTTP proxy will pre-fetch those embedded objects from the Web server without a request from the client. The objects are then ready to be delivered from the cache straight to the user, as fast as the client can request them.

While object pipelining enhances the user experience by minimizing latency and improving response times for first-time Web page requests, it could increase bandwidth utilization. Therefore by default, to avoid an increase in bandwidth utilization, object pipelining is disabled for the reverse proxy and bandwidth gain profiles. It is enabled, by default, only on the forward proxy — Normal profile, where enhancing the response time for clients is vital.

About HTTP Object Freshness

HTTP proxy categorizes HTTP objects into three types:

- ❑ Type-T: The OCS specifies explicit expiration time.
- ❑ Type-M: Expiration time is not specified; however, the last modified time is specified by the OCS.
- ❑ Type-N: Neither expiration nor last modified time has been specified.

The ProxySG Asynchronous Adaptive Refresh (AAR) algorithm was designed to maintain the freshness for all three types of cached HTTP objects in environments where the Internet was characterized by larger, static pages and relatively low Internet connection speeds. With AAR enabled, the ProxySG performs *freshness checks* with the OCS to expunge old content from cache and to replace it with updated content. To maximize the freshness of the next access to objects in the cache, the ProxySG appliance uses the AAR algorithm to perform asynchronous revalidations on those objects based on their relative popularity and the amount of time remaining before their estimated time of expiration.

However, with the advent of Web 2.0, the nature of the Internet has changed. With the general adoption of Web 2.0, which is characterized by dynamic content with many small objects coupled with increasing bandwidth to the Internet, the methods for caching are also evolving. In SGOS 6.2 the object cache model was changed to support more objects per disk to allow for better support for Web 2.0 content. With the addition of this object model, the value of the original adaptive refresh model has diminished markedly, and can in many instances actually increase the latency due to system load.

Therefore, AAR is now disabled by default on systems running SGOS 6.2.6 (and later). However, if you upgrade from a pre-SGOS 6.2.6 release, AAR may still be enabled. For information on how to configure this feature to best serve your environment, see ["Allocating Bandwidth to Refresh Objects in Cache"](#) on page 184.

About Meta Tags

A meta tag is a hidden tag that placed in the `<head>` of an HTML document. It provides descriptions and keywords for search engines and can contain the attributes — `content`, `http-equiv`, and `name`. Meta tags with an `http-equiv` attribute are equivalent to HTTP headers.

The ProxySG does not parse HTTP meta tag headers if:

- ❑ The meta tag does not appear within the first 256 bytes of the HTTP object body. To be parsed, relevant HTTP meta tags must appear within the first 256 bytes of the HTTP object body.
- ❑ The Blue Coat AV that is connected to your ProxySG, adds or modifies the meta tags in its response to the ProxySG. The response body modified by the Blue Coat AV is not parsed.

Planning Considerations

You can use CPL properties in the `<Cache>` layer to control meta tag processing. The CPL commands can be used in lieu of the check boxes for parsing meta tags through the Management Console. For details on the meta-tags, see Step 7 in "[To set HTTP default object caching policy](#):" on page 171.

The following CPL commands are applicable for HTTP proxy, HTTP refresh, and HTTP pipeline transactions:

```
http.response.parse_meta_tag.Cache-Control(yes|no)
http.response.parse_meta_tag.Expires(yes|no)
http.response.parse_meta_tag.Pragma.no-cache(yes|no)
```

VPM support to control the processing of meta tags is not available.

Related CLI Syntax to Parse Meta Tags

```
SGOS#(config) http [no] parse meta-tag cache-control
SGOS#(config) http [no] parse meta-tag expires
SGOS#(config) http [no] parse meta-tag pragma-no-cache
```

About Tolerant HTTP Request Parsing

The tolerant HTTP request parsing flag causes certain types of malformed requests to be processed instead of being rejected. The defaults are:

- ❑ Proxy Edition: The HTTP tolerant request parsing flag is not set, by default, The ProxySG blocks malformed HTTP requests, returning a *400 Invalid Request* error.
- ❑ MACH5 Edition: The HTTP tolerant request parsing flag is set by default. Malformed HTTP requests are not blocked.

Implementation of HTTP Tolerant Request Parsing

By default, a header line that does not begin with a `<Tab>` or space character must consist of a header name (which contains no `<Tab>` or space characters), followed by a colon and an optional value.

When the tolerant HTTP request parsing flag is either not set or is disabled, if the header name and required details are missing, the ProxySG blocks malformed HTTP requests and returns a *400 Invalid Request* error.

With tolerant request parsing enabled, a request header name is allowed to contain <Tab> or space characters, and if the request header line does not contain a colon, then the entire line is taken as the header name.

A header containing only one or more <Tab> or space characters is considered ambiguous. The ProxySG cannot discern if this is a blank continuation line or if it is a blank line that signals the end of the header section. By default, an ambiguous blank line is illegal, and an error is reported. With tolerant request parsing enabled, an ambiguous blank line is treated as the blank line that ends the header section.

To enable the HTTP tolerant request parsing flag:

Note: This feature is only available through the CLI.

From the (config) prompt, enter the following command to enable tolerant HTTP request parsing (the default is disabled):

```
SGOS#(config) http tolerant-request-parsing
```

To disable HTTP tolerant request parsing:

```
SGOS#(config) http no tolerant-request-parsing
```

About HTTP Compression

Compression reduces a file size but does not lose any data. Whether you should use compression depends upon three resources: server-side bandwidth, client-side bandwidth, and ProxySG CPU. If server-side bandwidth is more expensive in your environment than CPU, always request compressed content from the origin content server (OCS). However, if CPU is comparatively expensive, the ProxySG appliance should instead be configured to ask the OCS for the same compressions that the client requested and to forward whatever the server returns.

The default configuration assumes that CPU is costlier than bandwidth. If this is not the case, you can change the ProxySG appliance behavior.

Note: Decompression, content transformation, and recompression increases response time by a small amount because of the CPU overhead. (The overhead is negligible in most cases.) RAM usage also increases if compression is enabled.

Compression might also appear to adversely affect bandwidth gain. Because compression results in a smaller file being served to the client than was retrieved by the ProxySG appliance from the origin content server, bandwidth gain statistics reflect such requests/responses as negative bandwidth gain.

Compression is disabled by default. If compression is enabled, the HTTP proxy forwards the supported compression algorithm (gzip and deflate) from the client's request (Accept-Encoding: request header) to the server as is, and attempts to send compressed

content to client whenever possible. This allows the ProxySG appliance to send the response as is when the server sends compressed data, including non-cacheable responses. Any unsolicited encoded response is forwarded to the client as is.

Note: If compression is not enabled, the ProxySG appliance does not compress the content if the server sends uncompressed content. However, the appliance continues to uncompress content if necessary to apply transformations.

Any unsolicited encoded response is forwarded to the client as is.

Compression is controlled by policy only.

You can view compression statistics by going to **Statistics > Protocol Details > HTTP/FTP History > Client Comp. Gain** and **Server Comp. Gain**.

For information on these statistics, see "[Viewing HTTP/FTP Statistics](#)" on page 202.

Understand Compression Behavior

The ProxySG compression behavior is detailed in the tables below. Compression increases the overall percentage of cacheable content, increasing the hit rate in terms of number of objects served from the cache.

Note: A variant is the available form of the object in the cache—compressed or uncompressed. The Content-Encoding: header Identity refers to the uncompressed form of the content.

For cache-hit compression behavior, see [Table 8-1](#) below. For cache-miss compression behavior, see [Table 8-2](#).

Table 8-1. Cache-Hit Compression Behavior

Accept-Encoding: in client request	Variant Available when the Request Arrived	Variant Stored as a Result of the Request	Content-Encoding: in ProxySG response
Identity	Uncompressed object	None	Identity
Identity	No uncompressed object gzip compressed	Uncompressed	Identity
gzip, deflate	Uncompressed object	gzip compressed	gzip
gzip, deflate	Uncompressed object gzip compressed	None	gzip
gzip, deflate	Uncompressed object deflate compressed	None	deflate
deflate	No uncompressed object gzip compressed	deflate compressed	deflate (This is effectively a cache-miss. The ProxySG appliance does not convert from gzip to deflate.)

Table 8-2. Cache-Miss Compression Behavior

Accept-Encoding: in client request	Accept-Encoding: in ProxySG request	Content-Encoding: in server response	Generated variants	Content-Encoding: in ProxySG response
Identity	Identity	Identity	uncompressed object	Identity
gzip, deflate	gzip, deflate	Identity	uncompressed object gzip-compressed	gzip
gzip, deflate	gzip, deflate	gzip	No uncompressed object gzip-compressed	gzip
gzip, deflate, compress	gzip, deflate	gzip	No uncompressed object gzip-compressed	gzip
gzip, deflate	gzip, deflate	compress (illegal response)	compress	compress

Compression Exceptions

- ❑ The ProxySG appliance issues a `transformation_error` exception (HTTP response code 403), when the server sends an unknown encoding and the appliance is configured to do content transformation.
- ❑ The ProxySG appliance issues an `unsupported_encoding` exception (HTTP response code 415 - Unsupported Media Type) when the appliance is unable to deliver content due to configured policy.

The messages in the exception pages can be customized. For information on using exception pages, refer to “Advanced Policy Tasks” in the *Visual Policy Manager Reference*.

Configuring Compression

Compression behavior can only be configured through policy—VPM or CPL.

Using VPM to Configure Compression Behavior

Three objects can be used to configure compression and compression levels through VPM:

- ❑ Client HTTP compression object: Allows you to determine the behavior when the client wants the content in a different form than is in the cache.
- ❑ Server HTTP compression object: Allows you to enable or disable compression and to set options.
- ❑ HTTP compression level object: Allows you to set a compression level of low, medium, or high.

Refer to the *Visual Policy Manager Reference* to configure these HTTP compression options.

Using Policy to Configure Compression Behavior

Compression and decompression are allowed if compression is enabled. If compression is not enabled, neither compression nor decompression are allowed.

Policy controls the compression or decompression of content on the ProxySG appliance. If compression is turned off, uncompressed content is served to the client if a compressed variant is not available. If decompression is disabled, an uncompressed version is fetched from the OCS if the variant does not exist and the client requested uncompressed content.

Note: The ProxySG appliance decompresses the content if transformation is to be applied, even if the compression is not enabled.

You can use server-side or client-side controls to manage compression through policy, as described in the following table.

Table 8-3. Compression Properties

<i>Compression Properties</i>	<i>Description</i>
<code>http.allow_compression(yes no)</code>	Allow the ProxySG appliance to compress content on demand if needed.
<code>http.allow_decompression(yes no)</code>	Allow the ProxySG appliance to decompress content on demand if needed.
<code>http.compression_level(low medium high)</code>	Set the compression level to be low (1), medium (6), or high (9). Low is the default.
<code>http.server.accept_encoding(client)</code>	Turn on only client encodings
<code>http.server.accept_encoding(identity)</code>	Turn off all encodings
<code>http.server.accept_encoding(all)</code>	Turn on all supported encodings, including the client's encodings.
<code>http.server.accept_encoding(gzip, deflate)</code>	Send specific encodings (order sensitive)
<code>http.server.accept_encoding(gzip, client)</code>	Send specific encodings (order sensitive)
<code>http.server.accept_encoding.gzip(yes no)</code>	Add/remove an encoding
<code>http.server.accept_encoding[gzip, deflate, identity](yes no)</code>	Add/remove a list of encodings
<code>http.server.accept_encoding.allow_unknown(yes no)</code>	Allow/disallow unknown encodings.
<code>http.client.allow_encoding(identity);</code>	Allow no encodings (send uncompressed).
<code>http.client.allow_encoding(client);</code>	Allow all client encodings. This is the default.
<code>http.client.allow_encoding(gzip, deflate);</code>	Allow fixed set of encodings.

Table 8-3. Compression Properties (Continued)

Compression Properties	Description
<code>http.client.allow_encoding(gzip, client);</code>	Allow fixed set of encodings.
<code>http.client.allow_encoding.gzip(yes no);</code>	Add/remove one encoding
<code>http.client.allow_encoding[gzip, deflate, identity](yes no);</code>	Add/remove list of encodings

Default Behavior

By default, Blue Coat sends the client's list of the accept encoding algorithms, except for unknown encodings. If compression is not enabled, the default overrides any configured CPL policy.

If Accept-Encoding request header modification is used, it is overridden by the compression related policy settings shown in Table 8-3. The Accept-Encoding header modification can continue to be used if no compression policies are applied, or if compression is not enabled. Otherwise, the compression-related policies override any Accept-Encoding header modification, even if the Accept-Encoding header modification appears later in the policy file.

Adding encoding settings with client-side controls depend on if the client originally listed that encoding in its Accept-Encoding header. If so, these encodings are added to the list of candidates to be delivered to the client. The first cache object with an Accept-Encoding match to the client-side list is the one that is delivered.

Suggested Settings for Compression

- ❑ If client-side bandwidth is expensive in your environment, use the following policy:

```
<proxy>
  http.client.allow_encoding(client)
  http.allow_compression(yes)
```

- ❑ If server-side bandwidth is expensive in your environment, compared to client-side bandwidth and CPU:

```
http.server.accept_encoding(all)
http.server.accept_encoding.allow_unknown(no); default
http.allow_compression(yes)
http.allow_decompression(yes)
```

- ❑ If CPU is expensive in your environment, compared to server-side and client-side bandwidth:

```
http.server.accept_encoding(client); If no content transformation
policy is configured
http.server.accept_encoding(identity); If some content transformation
policy is configured
http.allow_compression(no); default
http.allow_decompression(no); default
```

Notes

- ❑ Policy-based content transformations are not stored as variant objects. If content transformation is configured, it is applied on all cache-hits, and objects might be compressed all the time at the end of such transformation if they are so configured.
- ❑ The variant that is available in the cache is served, even if the client requests a compression choice with a higher qvalue. For example, if a client requests `Accept-encoding: gzip;q=1, deflate;q=0.1`, and only a deflate-compressed object is available in the cache, the deflate compressed object is served.
- ❑ The HTTP proxy ignores `Cache-Control: no-transform` directive of the OCS. To change this, write policy to disallow compression or decompression if `Cache-Control: no-transform` response header is present.
- ❑ The ProxySG appliance treats multiple content encoding (gzip, deflate or gzip, gzip) as an unknown encoding. (These strings indicate the content has been compressed twice.)
- ❑ The gzip and deflate formats are treated as completely separate and are not converted from one to the other.
- ❑ Blue Coat recommends using gzip encoding (or allowing both gzip and deflate) when using the HTTP compression feature.
- ❑ If the ProxySG appliance receives unknown content encoding and if content transformation is configured (such as popup blocking), an error results.
- ❑ If the origin server provides compressed content with a different compression level than that specified in policy, the content is not re-compressed.
- ❑ If the ProxySG appliance compressed and cached content at a different compression level than the level specified in a later transaction, the content is not re-compressed.
- ❑ Parsing of container HTML pages occurs on the server side, so pipelining (prefetching) does not work when the server provides compressed content.
- ❑ Compressing a zip file breaks some browser versions, and compressing images does not provide added performance.
- ❑ All responses from the server can be compressed, but requests to the server, such as POST requests, cannot.
- ❑ Only 200 OK responses can be compressed.

About the HTTP Object Caching Policy Global Defaults

The ProxySG offers multiple configuration options that allow you to treat cached objects in a way that best suits your business model.

The following table lists the options that you can configure.

Table 8–1 Settings for Configuring the Object Caching Policy

Settings to Configure Object Caching	Notes
Setting the maximum object cache size	The default is 10000 MB for new installations of SGOS 6.5.9.10 and later. If you are using a version of SGOS prior to 6.5.2 or have upgraded to 6.5.2 or later, the default is 1024 MB.
Setting the TTL for negative responses in cache	Determines the number of minutes the SGOS stores negative responses for requests that could not be served to the client. The OCS might send a client error code (4xx response) or a server error code (5xx response) as a response to some requests. If you configure the ProxySG to cache negative responses for a specified number of minutes, it returns the negative response in subsequent requests for the same page or image for the specified length of time. The ProxySG will not attempt to fetch the request from the OCS. Therefore, while server-side bandwidth is saved, you could receive negative responses to requests that might otherwise have been served by accessing the OCS. By default, the ProxySG does not cache negative responses. It always attempts to retrieve the object from the OCS, if it is not already in cache. Default: 0 minutes
Forcing freshness validation before serving an object from cache	Verifies that each object is fresh upon access. Enabling this setting has a significant impact on performance because the HTTP proxy revalidates requested cached objects with the OCS before serving them to the client. This results in a negative impact on bandwidth gain. Therefore, do not enable this configuration unless absolutely required. For enabling, select the Always check with source before serving object check box. Default: Disabled
Parsing HTTP meta tag headers	Determines how HTTP meta tag headers are parsed in the HTML documents. The meta tags that can be enabled for parsing are: Cache-control meta tag The sub-headers that are parsed when this check box is selected are: <code>private</code>, <code>no-store</code>, <code>no-cache</code>, <code>max-age</code>, <code>s-maxage</code>, <code>must-revalidate</code>, <code>proxy-revalidate</code> Expires meta tag This directive parses for the date and time after which the document should be considered expired. Pragma-no-cache meta tag This directive indicates that cached information should not be used and instead requests should be forwarded to the OCS. Default: Disabled
Allocating bandwidth on the HTTP proxy for maintaining freshness of the objects in cache	Allows you to specify a limit to the amount of bandwidth the ProxySG uses to achieve the desired freshness. For more information see, "Allocating Bandwidth to Refresh Objects in Cache" on page 184. Default: Disable refreshing

The above settings serve as defaults on the proxy. If you want a more granular caching policy, for example— setting the TTL for an object, use Blue Coat Content Policy Language (CPL). You can also use the VPM or CPL to bypass the cache or to prohibit caching for a specific domain or server. Refer to the *Content Policy Language Guide* for more information.

About Clientless Requests Limits

When certain HTTP proxy configurations are enabled, the ProxySG employs various server-side connections to the OCS that are essential to caching and optimizing HTTP traffic. The ProxySG automatically sends requests, called clientless requests, over these connections. Performance and poor user experience might occur, however, when an unlimited number of clientless requests are allowed. As clientless requests increase and overwhelm the OCS, users might experience slow downloads in their Web browsers. Furthermore, these excessive requests might trigger the defensive measures because the corporate firewall determines that the ProxySG is a security threat.

The following sub-sections describe the HTTP proxy functionality involved.

HTTP Content Pre-population

Configuration: Symantec Director distributes content management commands; ProxySG connects to the OCS.

Symptom: The OCS becomes overwhelmed.




Figure 8–2 No Clientless Request Limits and HTTP Content Pre-population

The OCS becomes overwhelmed from content requests and content management commands. In this deployment, a global limit is not sufficient; a per-server limit is required.

Caching/Optimization (Pipelining)

Configuration: Blue Coat ProxySG pipelining options enabled (**Configuration > Proxy Settings > HTTP Proxy > Acceleration Profile**).

Symptom: The OCS becomes overwhelmed; users report slow access times in their Web browsers.




Figure 8–3 No Clientless Request Limits and Pipelining Enabled

Responses to clients might contain embedded links that the ProxySG converts to pipeline requests. As each link request results in a request to the OCS, performance might be impacted; if the firewall in front of the OCS determines that the request storm from the ProxySG represents a threat, requests are not allowed through. In this scenario, a per-page limit prevents the problem.

Bandwidth Gain

Configuration: Blue Coat ProxySG **Enable Bandwidth Gain Mode** option enabled (**Configuration > Proxy Settings > HTTP Proxy > Acceleration Profile**).

Symptom: The OCS becomes overwhelmed.



Figure 8–4 No Clientless Request Limits and Bandwidth Gain is Enabled

The ProxySG determines that objects in the cache require refreshing. This operation itself is not costly, but the additional requests to the OCS adds load to the WAN link. A global and per-server limit prevents the problem.

For new installations (or following a restoration to factory defaults), clientless limits are enforced by default; the ProxySG capacity per model determines the upper default limit. For systems upgraded to SGOS 6.x from versions previous to 5.x, clientless limits are not enforced and you must manually configure the ProxySG.

Continue with ["Setting the HTTP Default Object Caching Policy"](#) on page 171.

Preventing Exception Pages From Upstream Connection Errors

The ProxySG provides an option that prevents the ProxySG from returning TCP error exception pages to clients when upstream connection errors or connection time outs occur.

These types of connection issues might be common when enterprises employ custom applications. Though the connections issues are related to the server, administrators might mistakenly conclude that the ProxySG is the source of the problem because of the issues exception page from the proxy.

When the option is enabled, the ProxySG essentially closes connections to clients upon a server connection error or timeout. To the user, the experience is a lost connection, but not an indication that something between (such as a proxy) is at fault.

This feature is enabled (send exceptions on error) by default:

- ❑ After upgrading to SGOS 6.x from previous versions that have an Acceleration License

- On systems that have the acceleration profile selected during initial configuration (see [Section D: "Selecting an HTTP Proxy Acceleration Profile"](#) on page 173).

This option can only be enabled/disabled through the CLI:

```
SGOS#(config) http exception-on-network-error
SGOS#(config) http no exception-on-network-error
```

Setting the HTTP Default Object Caching Policy

This section describes how to set the HTTP default object caching policy. For more information, see ["HTTP Optimization"](#) on page 158.

To set HTTP default object caching policy:

- Verify that the ProxySG is intercepting HTTP traffic (**Configuration > Proxy Services; Standard** service group (by default)).
- From the Management Console, select **Configuration > Proxy Settings > HTTP Proxy > Policies**.

- Configure default proxy policies (**HTTP Proxy Policy** area; see ["About the HTTP Object Caching Policy Global Defaults"](#) on page 167):
 - In the **Do not cache objects larger than** field, enter the maximum object size to cache. The default is 10000 MB for new installations of SGOS 6.5.9.10 and later. If you are using a version of SGOS prior to 6.5.2 or have upgraded to 6.5.2 or later, the default is 1024 MB.
 - In the **Cache negative responses for** field, enter the number of minutes that SGOS stores negative responses. The default is 0.

- c. Force freshness validation. To always verify that each object is fresh upon access, select the **Always check with source before serving object** option. Enabling this setting has a significant impact on performance, do not enable this configuration unless absolutely required.
- d. Disable meta-tag parsing. The default is to parse HTTP meta tag headers in HTML documents if the MIME type of the object is text/html.

To disable meta-tag parsing, clear the option for:

- **Parse cache-control meta tag**
The following sub-headers are parsed when this check box is selected:
`private, no-store, no-cache, max-age, s-maxage, must-revalidate, proxy-revalidate.`
- **Parse expires meta tag**
This directive parses for the date and time after which the document should be considered expired.
- **Parse pragma-no-cache meta tag**
This directive indicates that cached information should not be used and instead requests should be forwarded to the OCS.

4. Configure **Clientless Request Limits** (see ["About Clientless Requests Limits"](#) on page 169):
 - a. **Global Limit**—Limits the number of concurrent clientless connections from the ProxySG to any OCS. Strongly recommended if **Pipeline** options or the **Enable Bandwidth Gain Mode** option is enabled on the **Configuration > Proxy Settings > HTTP Proxy > Acceleration Profile** tab.
 - b. **Per-server Limit**—Limits the number of concurrent clientless connections from the ProxySG to a specific OCS, as determined by the hostname of the OCS. Strongly recommended if **Pipeline** options or the **Enable Bandwidth Gain Mode** option is enabled on the **Configuration > Proxy Settings > HTTP Proxy > Acceleration Profile** tab.
 - c. **Per-page Limit**—Limits the number of requests that are created as a result of embedded objects.

5. Click **OK**; click **Apply**.

See Also

- ❑ ["Customizing the HTTP Object Caching Policy"](#) on page 158.
- ❑ ["Clearing the Object Cache"](#) on page 1403
- ❑ ["Selecting an HTTP Proxy Acceleration Profile"](#) on page 173.

Section D: Selecting an HTTP Proxy Acceleration Profile

This section discusses caching, pipelining behavior, and bandwidth gain.

Acceleration Profile Tasks

A proxy profile offers a collection of attributes that determine object caching and object pipelining behavior. The attributes are pre-selected to meet a specific objective — reduce response time for clients, reduce load on the OCS, reduce server-side bandwidth usage.

Based on your needs, you can select any of the three profiles offered or you can create a customized profile by selecting or clearing the options available within a profile.

The available proxy profile are:

- ❑ Normal (the default setting) acts as a client accelerator, and is used for enterprise deployments.
- ❑ Portal acts as a server accelerator (reverse proxy), and is used for Web hosting.
- ❑ Bandwidth Gain is used for Internet Service Provider (ISP) deployments.

Topic Links

- ❑ ["About the Normal Profile"](#)
- ❑ ["About the Portal Profile"](#)
- ❑ ["About the Bandwidth Gain Profile"](#) on page 174
- ❑ ["About HTTP Proxy Profile Configuration Components"](#)

About the Normal Profile

Normal is the default profile and can be used wherever the ProxySG is used as a normal forward proxy. This profile is typically used in enterprise environments, where the freshness of objects is more important than controlling the use of server-side bandwidth. The Normal profile is the profile that most follows the HTTP standards concerning object revalidation and staleness. Additionally, pre-fetching (pipelining) of embedded objects and redirects is enabled, which reduces response time for clients.

About the Portal Profile

When configured as a server accelerator or reverse proxy, the ProxySG improves object response time to client requests, scalability of the origin content server (OCS) site, and overall Web performance at the OCS. A server accelerator services requests meant for an OCS, as if it is the OCS itself.

About the Bandwidth Gain Profile

The Bandwidth Gain profile is useful wherever server-side bandwidth is an important resource. This profile is typically used in Internet Service Provider (ISP) deployments. In such deployments, minimizing server-side bandwidth is most important. Therefore, maintaining the freshness of an object in cache is less important than controlling the use of server-side bandwidth. The Bandwidth-Gain profile enables various HTTP configurations that can increase page response times and the likelihood that stale objects are served, but it reduces the amount of server-side bandwidth required.

About HTTP Proxy Profile Configuration Components

The following table describes each HTTP proxy acceleration profile option (Management Console and CLI).

Table 8–2 Description of Profile Configuration Components

Management Console Check box Field	CLI (config) Command	Definition
Pipeline embedded objects in client request	<code>http [no] pipeline client requests</code>	This configuration item applies only to HTML responses. When this setting is enabled, and the object associated with an embedded object reference in the HTML is not already cached, HTTP proxy acquires the object's content before the client requests the object. This improves response time dramatically. If this setting is disabled, HTTP proxy does not acquire embedded objects until the client requests them.
Pipeline redirects for client request	<code>http [no] pipeline client redirects</code>	When this setting is enabled, and the response of a client request is one of the redirection responses (such as 301, 302, or 307 HTTP response code), then HTTP proxy pipelines the object specified by the <code>Location</code> header of that response, provided that the redirection location is an HTML object. This feature improves response time for redirected URLs. If this setting is disabled, HTTP proxy does not pipeline redirect responses resulting from client requests.

Table 8–2 Description of Profile Configuration Components (Continued)

Management Console Check box Field	CLI (config) Command	Definition
Pipeline embedded objects in prefetch request	<code>http [no] pipeline prefetch requests</code>	This configuration item applies only to HTML responses resulting from pipelined objects. When this setting is enabled, and a pipelined object's content is also an HTML object, and that HTML object has embedded objects, then HTTP proxy also pipelines those embedded objects. This nested pipelining behavior can occur three levels deep at most. If this setting is disabled, the HTTP proxy does not perform nested pipelining.
Pipeline redirects for prefetch request	<code>http [no] pipeline prefetch redirects</code>	When this setting is enabled, HTTP proxy pipelines the object specified by a redirect location returned by a pipelined response. If this setting is disabled, HTTP proxy does not try to pipeline redirect locations resulting from a pipelined response.
Substitute Get for IMS	<code>http [no] substitute if- modified-since</code>	If the time specified by the <code>If-Modified-Since</code> header in the client's conditional request is greater than the last modified time of the object in the cache, it indicates that the copy in cache is stale. If so, HTTP proxy does a conditional GET to the OCS, based on the last modified time of the cached object. To change this aspect of the <code>If-Modified-Since</code> header on the ProxySG, enable the Substitute Get for IMS setting. When this setting is enabled, a client time condition greater than the last modified time of the object in the cache does not trigger revalidation of the object. Note: All objects do not have a last-modified time specified by the OCS.

Table 8–2 Description of Profile Configuration Components (Continued)

Management Console Check box Field	CLI (config) Command	Definition
Pipeline embedded objects in prefetch request	<code>http [no] pipeline prefetch requests</code>	This configuration item applies only to HTML responses resulting from pipelined objects. When this setting is enabled, and a pipelined object's content is also an HTML object, and that HTML object has embedded objects, then HTTP proxy also pipelines those embedded objects. This nested pipelining behavior can occur three levels deep at most. If this setting is disabled, the HTTP proxy does not perform nested pipelining.
Pipeline redirects for prefetch request	<code>http [no] pipeline prefetch redirects</code>	When this setting is enabled, HTTP proxy pipelines the object specified by a redirect location returned by a pipelined response. If this setting is disabled, HTTP proxy does not try to pipeline redirect locations resulting from a pipelined response.
Substitute Get for IMS	<code>http [no] substitute if- modified-since</code>	If the time specified by the <code>If-Modified-Since:</code> header in the client's conditional request is greater than the last modified time of the object in the cache, it indicates that the copy in cache is stale. If so, HTTP proxy does a conditional GET to the OCS, based on the last modified time of the cached object. To change this aspect of the <code>If-Modified-Since:</code> header on the ProxySG, enable the Substitute Get for IMS setting. When this setting is enabled, a client time condition greater than the last modified time of the object in the cache does not trigger revalidation of the object. Note: All objects do not have a last-modified time specified by the OCS.

Table 8–2 Description of Profile Configuration Components (Continued)

Management Console Check box Field	CLI (config) Command	Definition
Substitute Get for HTTP 1.1 conditionals	<code>http [no] substitute conditional</code>	HTTP 1.1 provides additional controls to the client over the behavior of caches concerning the staleness of the object. Depending on various <code>Cache-Control:</code> headers, the ProxySG can be forced to consult the OCS before serving the object from the cache. For more information about the behavior of various <code>Cache-Control:</code> header values, refer to RFC 2616. If the Substitute Get for HTTP 1.1 Conditionals setting is enabled, HTTP proxy ignores the following <code>Cache-Control:</code> conditions from the client request: • <code>"max-stale" ["=" delta-seconds]</code> • <code>"max-age" "=" delta-seconds</code> • <code>"min-fresh" "=" delta-seconds</code> • <code>"must-revalidate"</code> • <code>"proxy-revalidate"</code>
Substitute Get for PNC	<code>http [no] substitute pragma-no-cache</code>	Typically, if a client sends an HTTP GET request with a <code>Pragma: no-cache</code> or <code>Cache-Control: no-cache</code> header (for convenience, both are hereby referred to as PNC), a cache must consult the OCS before serving the content. This means that HTTP proxy always re-fetches the entire object from the OCS, even if the cached copy of the object is fresh. Because of this, PNC requests can degrade proxy performance and increase server-side bandwidth utilization. However, if the Substitute Get for PNC setting is enabled, then the PNC header from the client request is ignored (HTTP proxy treats the request as if the PNC header is not present at all).
Substitute Get for IE reload	<code>http [no] substitute ie-reload</code>	Some versions of Internet Explorer issue the <code>Accept: */*</code> header instead of the <code>Pragma: no-cache</code> header when you click Refresh. When an <code>Accept</code> header has only the <code>*/*</code> value, HTTP proxy treats it as a PNC header if it is a type-N object. You can control this behavior of HTTP proxy with the Substitute GET for IE Reload setting. When this setting is enabled, the HTTP proxy ignores the PNC interpretation of the <code>Accept: */*</code> header.

Table 8–2 Description of Profile Configuration Components (Continued)

Management Console Check box Field	CLI (config) Command	Definition
Never refresh before expiration	<code>http [no] strict-expiration refresh</code>	Applies only to cached type-T objects. For information on HTTP object types, see "About HTTP Object Freshness" on page 160. When this setting is enabled, SGOS does not asynchronously revalidate such objects before their specified expiration time. When this setting is disabled, such objects, if they have sufficient relative popularity, can be asynchronously revalidated and can, after a sufficient number of observations of changes, have their estimates of expiration time adjusted accordingly.
Never serve after expiration	<code>http [no] strict-expiration serve</code>	Applies only to cached type-T objects. If this setting is enabled, an object is synchronously revalidated before being served to a client, if the client accesses the object after its expiration time. If this setting is disabled, the object is served to the client and, depending on its relative popularity, may be asynchronously revalidated before it is accessed again.
Cache expired objects	<code>http [no] cache expired</code>	Applies only to type-T objects. When this setting is enabled, type-T objects that are already expired at the time of acquisition is cached (if all other conditions make the object cacheable). When this setting is disabled, already expired type-T objects become non-cacheable at the time of acquisition.

Table 8–2 Description of Profile Configuration Components (Continued)

Management Console Check box Field	CLI (config) Command	Definition
Enable Bandwidth Gain Mode	<code>bandwidth-gain {disable enable}</code>	This setting controls both HTTP-object acquisition after client-side abandonment and AAR (asynchronous adaptive refresh) revalidation frequency. • HTTP-Object Acquisition When Bandwidth Gain mode is enabled, if a client requesting a given object abandons its request, then HTTP proxy immediately abandons the acquisition of the object from the OCS, if such an acquisition is still in progress. When bandwidth gain mode is disabled, the HTTP proxy continues to acquire the object from the OCS for possible future requests for that object. • AAR Revalidation Frequency Under enabled bandwidth gain mode, objects that are asynchronously refreshable are revalidated at most twice during their estimated time of freshness. With bandwidth gain mode disabled, they are revalidated at most three times. Not all asynchronously refreshable objects are guaranteed to be revalidated.

When a ProxySG is first manufactured, it is set to a **Normal** profile. Depending on your needs, you can use the **Bandwidth Gain** profile or the **Portal** profile. You can also combine elements of all three profiles, as needed for your environment.

The following table provides the default configuration for each profile.

Table 8–3 Normal, Portal, and Bandwidth Gain Profiles

Configuration	Normal Profile	Portal Profile	Bandwidth Gain
Pipeline embedded objects in client requests	Enabled	Disabled	Disabled
Pipeline embedded objects in prefetch requests	Enabled	Disabled	Disabled
Pipeline redirects for client requests	Enabled	Disabled	Disabled
Pipeline redirects for prefetch requests	Enabled	Disabled	Disabled
Cache expired objects	Enabled	Disabled	Enabled
Bandwidth Gain Mode	Disabled	Disabled	Enabled
Substitute GET for IMS (if modified since)	Disabled	Enabled	Enabled
Substitute GET for PNC (Pragma no cache)	Disabled	Enabled	Disabled

Table 8–3 Normal, Portal, and Bandwidth Gain Profiles (Continued)

Configuration	Normal Profile	Portal Profile	Bandwidth Gain
Substitute GET for HTTP 1.1 conditionals	Disabled	Enabled	Enabled
Substitute GET for IE (Internet Explorer) reload	Disabled	Enabled	Disabled
Never refresh before expiration	Disabled	Enabled	Enabled
Never serve after expiration	Enabled	Enabled	Disabled


Configuring the HTTP Proxy Profile

Configure the profile by selecting any of the components discussed in ["About HTTP Proxy Profile Configuration Components"](#) on page 174.

To configure the HTTP proxy profile:

1. Review the description of the components for each profile, see [Table 8–2](#) on page 174.
2. From the Management Console, select **Configuration > Proxy Settings > HTTP Proxy > Acceleration Profile**.

Text displays at the bottom of this tab indicating which profile is selected. Normal is the default profile. If you have a customized profile, this text does not display.



Important: If you have a customized profile and you click one of the **Use Profile** buttons, no record of your customized settings remains. However, after the ProxySG is set to a specific profile, the profile is maintained in the event the ProxySG is upgraded.

Also, if you select any **Pipeline** option or the **Enable Bandwidth Gain Mode** option, Symantec strongly recommends limiting clientless requests. See ["About Clientless Requests Limits"](#) on page 169.

3. To select a profile, click one of the three profile buttons (**Use Normal Profile**, **Use Bandwidth Gain Profile**, or **Use Portal Profile**).

The text at the bottom of the **Acceleration Profile** tab changes to reflect the new profile.

Note: You can customize the settings, no matter which profile button you select.

4. (Optional) To customize the profile settings, select or clear any of the check boxes (see [Table 8–2, "Description of Profile Configuration Components"](#) on page 174 for information about each setting).
5. Click **OK**; click **Apply**.

See Also

- ❑ ["Selecting an HTTP Proxy Acceleration Profile"](#) on page 173.
- ❑ ["About HTTP Proxy Profile Configuration Components"](#) on page 174.
- ❑ ["About HTTP Object Freshness"](#) on page 160.
- ❑ ["Using a Caching Service"](#) on page 182.

Section E: Using a Caching Service

CachePulse is a caching service that provides you with optimal bandwidth gains for popular or high-bandwidth websites. Utilizing highly effective Web caching technology, CachePulse saves bandwidth on expensive international links and backhaul traffic, thereby improving Web experience for users.

CachePulse accelerates the delivery of rich Web 2.0 content, video, and large files such as:

- YouTube videos
- Netflix streaming media
- Microsoft Windows updates

Subscribing to the CachePulse service eliminates the need to maintain caching policy; when you first enable the service, it downloads the latest version of the caching policy database. CachePulse periodically updates the database as long as the service is enabled and an Internet connection exists.

Prerequisite for Using CachePulse

Before you can use CachePulse, you must have a valid license for the feature. Refer to your Sales Engineer for more information.

If you do not have a valid license, the Management Console might display Health Monitoring errors. The event log might also contain error messages about the subscription.

Enabling CachePulse

To enable CachePulse:

1. In the Management Console, select **Configuration > Proxy Settings > General**.
2. In the CachePulse section, select **Enable**.
3. Click **Apply**.

The appliance attempts to download the database.

What if the Initial Download is Not Successful?

If you receive a download error and the Management Console banner displays **Critical** shortly after you click **Apply**, the CachePulse database download might have failed. Check your network configuration and make sure that the appliance can connect to the Internet. Because the appliance attempts to communicate with the Symantec server over a secured connection on port 443, you might also have to allow outbound connections from the appliance on port 443 in the firewall.

To check if there was a download problem, select **Statistics > Health Monitoring > Status** and look for the status “CachePulse failed on initial download” for **Subscription Communication Status**.

See Also

- ❑ "Downloading the CachePulse Database"
- ❑ "About the Status Metrics" on page 1343

Downloading the CachePulse Database

You can download the CachePulse database at any time if the feature is enabled. If the initial download failed, and you resolved the issue that caused the failure, you can use this method to download database updates.

To download the CachePulse database:

1. In the Management Console, select **Configuration > Proxy Settings > General**.
2. In the CachePulse section, click **Download Now**.

The License and Download Status field shows statistics about the previous successful and unsuccessful downloads. If the last download was unsuccessful, the field contains an error.

General

Tunnel on Protocol Error

☐ TCP tunnel requests when a protocol error is detected

Reflect Client IP

☐ Reflect client's source IP when connecting to servers

Trust Destination IP

☐ Trust client-provided destination IP when connecting to servers

User Overflow Action

☒ Do not enforce licensed user limit

☐ Bypass connections from users over licensed limit

☐ Queue connections from users over licensed limit

CachePulse

☒ Enable

License and Download Status:

Download method: Direct

Unsuccessful downloads: 1

Last attempt:

Time: Mon, 29 Jul 2013 17:30:53 UTC

Downloading from: https://subscription.es.bluecoat.com/cachepulse/policy

% Failed to DNS resolve subscription.es.bluecoat.com

Last successful download: N/A

Download Now

Preview Apply Revert Help

If you receive a download error, check your network configuration and make sure that the appliance can connect to the Internet.

Section F: Fine-Tuning Bandwidth Gain

In addition to the components related to top-level profiles, other configurable items affect bandwidth gain. You can set the top-level profile (see ["Selecting an HTTP Proxy Acceleration Profile"](#) on page 173) and adjust the following configuration items to fine tune the ProxySG for your environment:

- ❑ Allocating bandwidth to refresh objects in cache
- ❑ Using Byte-range support
- ❑ Enabling the Revalidate pragma-no-cache (PNC)


Allocating Bandwidth to Refresh Objects in Cache

The **Refresh bandwidth** options control the server-side bandwidth used for all forms of asynchronous adaptive refresh activity. On systems with increased object store capacity, the value of asynchronous adaptive refresh has diminished markedly, and can in many instances actually increase latency due to system load. Therefore, this feature is disabled by default. You can select from the following options:

- ❑ **Disable refreshing**—Disables adaptive refresh. This setting is recommended on systems that use an increased object capacity disk model (SGOS 6.2 and later). This is the default setting for fresh installations of SGOS 6.2.6 and later.
- ❑ **Let the SG appliance manage refresh bandwidth**—The appliance will automatically use whatever bandwidth is available in its efforts to maintain 99.9% estimated freshness of the next access. You can also enable this from the CLI using the `SGOS#(config caching) refresh bandwidth automatic` command. This setting is recommended only on systems that are not using the increased object capacity disk model (that is, systems that were manufactured with an SGOS version prior to 6.2).
- ❑ **Limit refresh bandwidth to x kilobits/sec**—If you want to use adaptive refresh but you want to limit the amount of bandwidth used, select this option and specify a limit to the amount of bandwidth the ProxySG uses to achieve the desired freshness. Before making adjustments, review the logged statistics and examine the current bandwidth used as displayed in the **Refresh bandwidth** field. It is not unusual for bandwidth usage to spike occasionally, depending on access patterns at the time. Entering a value of zero disables adaptive refresh.

To set refresh bandwidth:

1. From the Management Console, select **Configuration > Proxy Settings > HTTP Proxy > Freshness**.



The **Refresh bandwidth** field displays the refresh bandwidth options. The default setting is to **Disable refreshing**.

Important: Blue Coat strongly recommends that you not change the setting from the default if you have a system with an increased object store capacity.

2. To enable adaptive refresh, select one of the following options:
 - Select **Limit refresh bandwidth to** and enter a bandwidth limit to use in the **kilobits/sec** field.
 - To allow the appliance to automatically determine the amount of bandwidth to use for adaptive refresh, select **Let the SG Appliance manage refresh bandwidth (recommended)**.
3. Click **OK**; click **Apply**.

Using Byte-Range Support

Byte-range support is an HTTP feature that allows a client to use the `Range: HTTP` header for requesting a portion of an object rather than the whole object. The HTTP proxy supports byte-range support and it is enabled by default.

When Byte-Range Support is Disabled

If byte-range support is disabled, HTTP treats all byte-range requests as non-cacheable. Such requests are never served from the cache, even if the object exists in the cache. The client's request is sent unaltered to the OCS and the response is not cached. Thus, a byte-range request has no effect on the cache if byte-range support is disabled.

When Byte-Range Support is Enabled

If the object is already in cache, the ProxySG serves the byte-range request from the cache itself. However, if the client's request contains a PNC header, the ProxySG always bypasses the cache and serves the request from the OCS.

If the object is not in cache, the ProxySG always attempts to minimize delay for the client.

- ❑ If the byte-range requested is near the beginning of the object, that is the start byte of the request is within 0 to 14336 bytes, then the ProxySG fetches the entire object from the OCS and caches it. However, the client is served the requested byte-range only.
- ❑ If the byte-range requested is not near the beginning of the object, that is the start byte of the request is greater than 14336 bytes, then the ProxySG fetches only the requested byte-range from the OCS, and serves it to the client. The response is not cached.

Note: The HTTP proxy never caches partial objects, even if byte-range support is enabled.

Since the ProxySG never caches partial objects, bandwidth gain is significantly affected when byte-range requests are used heavily. If, for example, several clients request an object where the start byte offset is greater than 14336 bytes, the object is never cached. The ProxySG fetches the same object from the OCS for each client, thereby causing negative bandwidth gain.

Further, download managers like NetAnts® typically use byte-range requests with PNC headers. To improve bandwidth gain by serving such requests from cache, enable the **revalidate pragma-no-cache** option along with byte-range support. See ["Enabling Revalidate Pragma-No-Cache"](#) on page 187.

To configure byte-range support:

Note: Enabling or disabling byte-range support can only be configured through the CLI.

To enable or disable byte-range support, enter one of the following commands at the (config) command prompt:

```
SGOS#(config) http byte-ranges
-or-
SGOS#(config) http no byte-ranges
```

Enabling Revalidate Pragma-No-Cache

The pragma-no-cache (PNC) header in a client's request causes the HTTP proxy to re-fetch the entire object from the OCS, even if the cached copy of the object is fresh. This roundtrip for PNC requests can degrade proxy performance and increase server-side bandwidth utilization.

While the **Substitute Get for PNC** configuration completely ignores PNC in client requests and potentially serves stale content, the `revalidate-pragma-no-cache` setting allows you to selectively implement PNC.

When the `revalidate-pragma-no-cache` setting is enabled, a client's non-conditional PNC-GET request results in a conditional GET request sent to the OCS if the object is already in cache. The `revalidate-pragma-no-cache` request allows the OCS to return the 304 Not Modified response, if the content in cache is still fresh. Thereby, the server-side bandwidth consumed is lesser as the full content is not retrieved again from the OCS.

By default, the revalidate PNC configuration is disabled and is not affected by changes in the top-level profile. When the **Substitute Get for PNC** configuration is enabled (see [Table 8–2, "Description of Profile Configuration Components"](#) on page 174 for details), the revalidate PNC configuration has no effect.

To configure the revalidate PNC setting:

Note: The `revalidate pragma-no-cache` setting can only be configured through the CLI.

To enable or disable the revalidate PNC setting, enter one of the following commands at the (config) command prompt:

```
SGOS#(config) http revalidate-pragma-no-cache
-or-
SGOS#(config) http no revalidate-pragma-no-cache
```

Interpreting Negative Bandwidth Gain Statistics

Bandwidth gain represents the overall bandwidth benefit achieved by object and byte caching, compression, protocol optimization, and object caching.

Occasionally, you might notice negative bandwidth gain when using the bandwidth gain profile. This negative bandwidth gain is observed because the

client-side cumulative bytes of traffic is lower than the server-side cumulative bytes of traffic for a given period of time. It is represented as a unit-less multiplication factor and is computed by the ratio:

$$\text{client bytes} / \text{server bytes}$$

Some factors that contribute to negative bandwidth gain are:

❑ Abandoned downloads (`delete_on_abandonment (no)`)

When a client cancels a download, the ProxySG continues to download the requested file to cache it for future requests. Since the client has cancelled the download, server-side traffic persists while the client-side traffic is halted. This continued flow of traffic on the server-side causes negative bandwidth gain.

Further with (`delete_on_abandonment (yes)`), when a client cancels a download, the ProxySG terminates the connection and stops sending traffic to the client. However, the server may have sent additional traffic to the ProxySG before it received the `TCP RESET` from the ProxySG. This surplus also causes negative bandwidth gain.

❑ Refreshing of the cache

Bandwidth used to refresh contents in the cache contributes to server-side traffic. Since this traffic is not sent to the client until requested, it might cause negative bandwidth gain.

❑ Byte-range downloads

When download managers use an open-ended byte-range, such as `Range: bytes 10000-`, and reset the connection after downloading the requested byte-range. The packets received by the ProxySG from the server are greater than those served to the client, causing negative bandwidth gain.

❑ Download of uncompressed content

If the ProxySG downloads uncompressed content, but compresses it before serving the content to the client, server-side traffic will be greater than client-side traffic. This scenario is typical in a reverse proxy deployment, where the server offloads the task of gzipping the content to the ProxySG.

❑ Reduced client-side throughput

In the short term, you will notice negative bandwidth gain if the client-side throughput is lower than the server-side throughput. If, for example, the ProxySG takes five minutes to download a 100 Mb file and takes 10 minutes to serve the file to the client. The ProxySG reflects negative bandwidth gain for the first five minutes.

To view bandwidth usage and bandwidth gain statistics on the HTTP proxy, click **Statistics > Traffic History** tab. Select the HTTP proxy service to view statistics over the last hour, day, week, month, and year. See ["Statistics"](#) on page 667 for information on the graphs.

Compression

Compression is disabled by default. If compression is enabled, the HTTP proxy forwards the supported compression algorithm (either deflate or gzip) from the client's request (`Accept-Encoding`: request header) to the server as is, and attempts to send compressed content to client whenever possible. This allows SGOS to send the response as is when the server sends compressed data, including non-cacheable responses. Any unsolicited encoded response is forwarded as is to the client.

For more information on compression, see ["Understanding HTTP Compression"](#) on page 193.

Related CLI Syntax to Configure HTTP

The following commands allow you to manage settings for an HTTP proxy.

Use the command below to enter the configuration mode.

```
SGOS# conf t
```

The following subcommands are available:

```
SGOS#(config) http [no] add-header client-ip
SGOS#(config) http [no] add-header front-end-https
SGOS#(config) http [no] add-header via
SGOS#(config) http [no] add-header x-forwarded-for
SGOS#(config) http [no] byte-ranges
SGOS#(config) http [no] cache authenticated-data
SGOS#(config) http [no] cache expired
SGOS#(config) http [no] cache personal-pages
SGOS#(config) http [no] force-ntlm
SGOS#(config) http ftp-proxy-url root-dir
SGOS#(config) http ftp-proxy-url user-dir
SGOS#(config) http [no] parse meta-tag {cache-control | expires |
pragma-no-cache}
SGOS#(config) http [no] persistent client
SGOS#(config) http [no] persistent server
SGOS#(config) http [no] persistent-timeout client num_seconds
SGOS#(config) http [no] persistent-timeout server num_seconds
SGOS#(config) http [no] pipeline client {requests | redirects}
SGOS#(config) http [no] pipeline prefetch {requests | redirects}
SGOS#(config) http [no] proprietary-headers bluecoat
SGOS#(config) http receive-timeout client num_seconds
SGOS#(config) http receive-timeout refresh num_seconds
SGOS#(config) http receive-timeout server num_seconds
SGOS#(config) http [no] revalidate-pragma-no-cache
SGOS#(config) http [no] strict-expiration refresh
SGOS#(config) http [no] strict-expiration serve
SGOS#(config) http [no] strip-from-header
SGOS#(config) http [no] substitute conditional
SGOS#(config) http [no] substitute ie-reload
SGOS#(config) http [no] substitute if-modified-since
SGOS#(config) http [no] substitute pragma-no-cache
SGOS#(config) http [no] tolerant-request-parsing
```

```
SGOS#(config) http upload-with-pasv disable
SGOS#(config) http upload-with-pasv enable
SGOS#(config) http version {1.0 | 1.1}
SGOS#(config) http [no] www-redirect
SGOS#(config) http [no] xp-rewrite-redirect
```

Note: For detailed information about using these commands, refer to the *Command Line Interface Reference*.

Section G: Caching Authenticated Data (CAD) and Caching Proxy Authenticated Data (CPAD)

This section describes how the ProxySG caches authenticated content over HTTP. Authentication over HTTP allows a user to prove their identity to a server or an upstream proxy to gain access to a resource.

The ProxySG uses CAD and CPAD to facilitate object caching at the edge and to help validate user credentials. Object caching in the ProxySG allows for lesser bandwidth usage and faster response times between the client and the server or proxy.

The deployment of the ProxySG determines whether it performs CAD or CPAD:

- ❑ When the Origin Content Server (OCS) performs authentication, the ProxySG performs CAD.
- ❑ When the upstream HTTP Proxy performs authentication, the downstream HTTP proxy or ProxySG executes CPAD.

About Caching Authenticated Data (CAD)

In the CAD scenario, when a user requests a resource that needs authentication, the OCS sends an `HTTP 401` error response to the user. The `HTTP 401` response also contains information on the authentication schemes that the OCS supports. To prove their identity to the OCS, the user resubmits the initial request along with the authentication details.




Figure 8-5 CAD: 200 response from the Origin Content Server.

The OCS then sends back one of the following responses:

- ❑ `HTTP 200` response status, authentication is accepted. The user receives the requested resource.

- HTTP 403 response status, user is not allowed to view the requested resource. The user is authenticated but is not authorized to receive the content, hence the user receives an error message.

When another user accesses the same URL, the ProxySG authenticates the user with the OCS and verifies the freshness of the content using the `Get If Modified Since` request. If the user is authorized and the content has not been modified, the OCS returns an HTTP 304 response message to the ProxySG. The ProxySG then serves the content from cache.

If the content has been modified, the OCS returns the HTTP 200 response along with the modified content.




Figure 8-6 CAD: 403 and 304 response codes from the OCS

Note: CAD is applicable only for pure HTTP authentication — the ProxySG caches authenticated data only when the OCS includes the `www-Authenticate` response code in the 401 response header. If, for example, the client accesses an OCS that uses forms-based authentication, the ProxySG does not perform CAD.

About Caching Proxy Authenticated Data (CPAD)

The CPAD deployment uses two ProxySG appliances — a local proxy and a gateway proxy. Figure 8-7 on page 193 below depicts the ProxySG appliances in a CPAD deployment.

When the user requests a resource, ProxySG1 forwards the request to ProxySG2. ProxySG2 issues the authentication challenge back to the user (a 407 response instead of the 401 response that the OCS serves). Upon successful authentication, ProxySG2 forwards the request to the OCS and the resource is served to the user.




Figure 8-7 CPAD: 200 response from ProxySG 2

In Figure 8-8, ProxySG1 caches proxy authenticated data and ProxySG2 performs authentication (instead of the OCS).




Figure 8-8 CPAD: 407 and 304 responses in a CPAD deployment

For subsequent users who access the same URL, see Figure 8-4, ProxySG1 forwards all requests to ProxySG2 with the `Get If Modified Since` request.

ProxySG2 issues the authentication challenge and provides one of the following responses:

- ❑ HTTP 200 response status, the user is allowed access to the requested resource but the content has changed.
- ❑ HTTP 304 response status, the user is authorized and the content can be served from the cache.
- ❑ HTTP 403 response status, the user is not authorized to view the requested resource.
- ❑ HTTP 407 response status, the user provided invalid credentials.

Understanding HTTP Compression

Compression reduces a file size but does not lose any data. Whether you should use compression depends upon three resources: server-side bandwidth, client-side bandwidth, and ProxySG CPU. If server-side bandwidth is more expensive

in your environment than CPU, always request compressed content from the origin content server (OCS). However, if CPU is comparatively expensive, the ProxySG appliance should instead be configured to ask the OCS for the same compressions that the client asked for and to forward whatever the server returns.

The default configuration assumes that CPU is costlier than bandwidth. If this is not the case, you can change the ProxySG appliance behavior.

Note: Decompression, content transformation, and recompression increases response time by a small amount because of the CPU overhead. (The overhead is negligible in most cases.) RAM usage also increases if compression is enabled.

Compression might also appear to adversely affect bandwidth gain. Because compression results in a smaller file being served to the client than was retrieved by the ProxySG appliance from the origin content server, bandwidth gain statistics reflect such requests/responses as negative bandwidth gain.

Compression is disabled by default. If compression is enabled, the HTTP proxy forwards the supported compression algorithm (gzip and deflate) from the client's request (`Accept-Encoding`: request header) to the server as is, and attempts to send compressed content to client whenever possible. This allows the ProxySG appliance to send the response as is when the server sends compressed data, including non-cacheable responses. Any unsolicited encoded response is forwarded to the client as is.

Note: If compression is not enabled, the ProxySG appliance does not compress the content if the server sends uncompressed content. However, the appliance continues to uncompress content if necessary to apply transformations.

Any unsolicited encoded response is forwarded to the client as is.

Compression is controlled by policy only.

You can view compression statistics by going to **Statistics > Protocol Details > HTTP/FTP History > Client Comp. Gain** and **Server Comp. Gain**.

For information on these statistics, see "[Viewing HTTP/FTP Statistics](#)" on page 202.

Understand Compression Behavior

The ProxySG compression behavior is detailed in the tables below. Compression increases the overall percentage of cacheable content, increasing the hit rate in terms of number of objects served from the cache.

Note: A variant is the available form of the object in the cache—compressed or uncompressed. The Content-Encoding: header Identity refers to the uncompressed form of the content.

For cache-hit compression behavior, see [Table 8-3](#) below. For cache-miss compression behavior, see [Table 8-4](#).

Table 8–4 Cache-Hit Compression Behavior

Accept-Encoding: in client request	Variant Available when the Request Arrived	Variant Stored as a Result of the Request	Content-Encoding: in ProxySG response
Identity	Uncompressed object	None	Identity
Identity	No uncompressed object gzip compressed	Uncompressed	Identity
gzip, deflate	Uncompressed object	gzip compressed	gzip
gzip, deflate	Uncompressed object gzip compressed	None	gzip
gzip, deflate	Uncompressed object deflate compressed	None	deflate
deflate	No uncompressed object gzip compressed	deflate compressed	deflate (This is effectively a cache-miss. The ProxySG appliance does not convert from gzip to deflate.)

Table 8–5 Cache-Miss Compression Behavior

Accept-Encoding: in client request	Accept-Encoding: in ProxySG request	Content-Encoding: in server response	Generated variants	Content-Encoding: in ProxySG response
Identity	Identity	Identity	uncompressed object	Identity
gzip, deflate	gzip, deflate	Identity	uncompressed object gzip-compressed	gzip
gzip, deflate	gzip, deflate	gzip	No uncompressed object gzip-compressed	gzip
gzip, deflate, compress	gzip, deflate	gzip	No uncompressed object gzip-compressed	gzip
gzip, deflate	gzip, deflate	compress (illegal response)	compress	compress

Compression Exceptions

- ❑ The ProxySG appliance issues a `transformation_error` exception (HTTP response code 403), when the server sends an unknown encoding and the appliance is configured to do content transformation.
- ❑ The ProxySG appliance issues an `unsupported_encoding` exception (HTTP response code 415 - Unsupported Media Type) when the appliance is unable to deliver content due to configured policy.

The messages in the exception pages can be customized. For information on using exception pages, The messages in the exception pages can be customized. For information on using exception pages, refer to the Advanced Policy Tasks chapter, Section E, of the *Visual Policy Manager Reference*.

Configuring Compression

Compression behavior can only be configured through policy—VPM or CPL.

Using VPM to Configure Compression Behavior


Three objects can be used to configure compression and compression levels through VPM:

- ❑ Client HTTP compression object: Allows you to determine the behavior when the client wants the content in a different form than is in the cache.
- ❑ Server HTTP compression object: Allows you to enable or disable compression and to set options.
- ❑ HTTP compression level object: Allows you to set a compression level of low, medium, or high.

Complete the following steps to manage server and client HTTP compression and compression levels.

To add or edit client compression:

1. Create a Web Access Layer:
 - a. From the Management Console, select **Configuration > Policy > Visual Policy Manager**; click **Launch**.
 - b. Select **Policy > Add Web Access Layer** from the menu of the Blue Coat VPM window that appears.
 - c. Type a layer name into the dialog that appears and click **OK**.
2. Add an Action object:
 - a. Right click on the item in the **Action** column; select **Set**.
 - b. Click **New** in the Set Action Object dialog that appears; select Set **Client HTTP Compression**.



- c. Select the compression options you want to use; click **OK**.
- d. Click **OK** again; close the VPM window and click **Yes** in the dialog to save your changes.

To add or edit server compression:

1. Create a Web Access Layer:
 - a. From the Management Console, select **Configuration > Policy > Visual Policy Manager**; click **Launch**.
 - b. Select **Policy > Add Web Access Layer** from the menu of the Blue Coat VPM window that appears.
 - c. Type a layer name into the dialog that appears and click **OK**.
2. Add an Action object:
 - a. Right click on the item in the **Action** column; select **Set**.
 - b. Click **New** in the Set Action Object dialog that appears; select **Set Server HTTP Compression**.
 - c. Select compression options; click **OK**.
 - d. Click **OK** again; close the VPM window and click **Yes** in the dialog to save your changes.

Using VPM to Set HTTP Compression Levels

You can control the compression level based on any transaction condition (such as the client IP address, the hostname, request/response headers, and the like).

To set compression levels:

1. Create a Web Access Layer:
 - From the Management Console, select **Configuration > Policy > Visual Policy Manager**; click **Launch**.

- Select **Policy > Add Web Access Layer** from the menu of the Blue Coat VPM window that appears.
 - Type a layer name into the dialog that appears and click **OK**.
2. Add an Action object:
- Right click on the item in the **Action** column; select **Set**.
 - Click **New** in the Set Action Object dialog that appears; select Set **HTTP Compression Level**.
 - Select the compression level needed; click **OK**.
 - Click **OK** again; close the VPM window and click **Yes** in the dialog to save your changes.

Using Policy to Configure Compression Behavior

Compression and decompression are allowed if compression is enabled. If compression is not enabled, neither compression nor decompression are allowed.

Policy controls the compression or decompression of content on the ProxySG appliance. If compression is turned off, uncompressed content is served to the client if a compressed variant is not available. If decompression is disabled, an uncompressed version is fetched from the OCS if the variant does not exist and the client requested uncompressed content.

Note: The ProxySG appliance decompresses the content if transformation is to be applied, even if the compression is not enabled.

You can use server-side or client-side controls to manage compression through policy, as described in the following table.

Table 8–6 Compression Properties

Compression Properties	Description
<code>http.allow_compression(yes no)</code>	Allow the ProxySG appliance to compress content on demand if needed.
<code>http.allow_decompression(yes no)</code>	Allow the ProxySG appliance to decompress content on demand if needed.
<code>http.compression_level(low medium high)</code>	Set the compression level to be low (1), medium (6), or high (9). Low is the default.
<code>http.server.accept_encoding(client)</code>	Turn on only client encodings
<code>http.server.accept_encoding(identity)</code>	Turn off all encodings
<code>http.server.accept_encoding(all)</code>	Turn on all supported encodings, including the client's encodings.
<code>http.server.accept_encoding(gzip, deflate)</code>	Send specific encodings (order sensitive)

Table 8–6 Compression Properties (Continued)

Compression Properties	Description
<code>http.server.accept_encoding(gzip, client)</code>	Send specific encodings (order sensitive)
<code>http.server.accept_encoding.gzip(yes no)</code>	Add/remove an encoding
<code>http.server.accept_encoding[gzip, deflate, identity](yes no)</code>	Add/remove a list of encodings
<code>http.server.accept_encoding.allow_unknown(yes no)</code>	Allow/disallow unknown encodings.
<code>http.client.allow_encoding(identity);</code>	Allow no encodings (send uncompressed).
<code>http.client.allow_encoding(client);</code>	Allow all client encodings. This is the default.
<code>http.client.allow_encoding(gzip, deflate);</code>	Allow fixed set of encodings.
<code>http.client.allow_encoding(gzip, client);</code>	Allow fixed set of encodings.
<code>http.client.allow_encoding.gzip(yes no);</code>	Add/remove one encoding
<code>http.client.allow_encoding[gzip, deflate, identity](yes no);</code>	Add/remove list of encodings

Default Behavior

By default, Blue Coat sends the client's list of the accept encoding algorithms, except for unknown encodings. If compression is not enabled, the default overrides any configured CPL policy.

If `Accept-Encoding` request header modification is used, it is overridden by the compression related policy settings shown in [Table 8-5](#). The `Accept-Encoding` header modification can continue to be used if no compression policies are applied, or if compression is not enabled. Otherwise, the compression-related policies override any `Accept-Encoding` header modification, even if the `Accept-Encoding` header modification appears later in the policy file.

Adding encoding settings with client-side controls depend on if the client originally listed that encoding in its `Accept-Encoding` header. If so, these encodings are added to the list of candidates to be delivered to the client. The first cache object with an `Accept-Encoding` match to the client-side list is the one that is delivered.

Suggested Settings for Compression

- ❑ If client-side bandwidth is expensive in your environment, use the following policy:

```
<proxy>
  http.client.allow_encoding(client)
  http.allow_compression(yes)
```

- ❑ If server-side bandwidth is expensive in your environment, compared to client-side bandwidth and CPU:

```
http.server.accept_encoding(all)
http.server.accept_encoding.allow_unknown(no); default
http.allow_compression(yes)
http.allow_decompression(yes)
```

- ❑ If CPU is expensive in your environment, compared to server-side and client-side bandwidth:

```
http.server.accept_encoding(client);If no content transformation
policy is configured
http.server.accept_encoding(identity);If some content transformation
policy is configured
http.allow_compression(no); default
http.allow_decompression(no); default
```

Notes

- ❑ Policy-based content transformations are not stored as variant objects. If content transformation is configured, it is applied on all cache-hits, and objects might be compressed all the time at the end of such transformation if they are so configured.
- ❑ The variant that is available in the cache is served, even if the client requests a compression choice with a higher qvalue. For example, if a client requests `Accept-encoding: gzip;q=1, deflate;q=0.1`, and only a deflate-compressed object is available in the cache, the deflate compressed object is served.
- ❑ The HTTP proxy ignores `Cache-Control: no-transform` directive of the OCS. To change this, write policy to disallow compression or decompression if `Cache-Control: no-transform` response header is present.
- ❑ The ProxySG appliance treats multiple content encoding (gzip, deflate or gzip, gzip) as an unknown encoding. (These strings indicate the content has been compressed twice.)
- ❑ The gzip and deflate formats are treated as completely separate and are not converted from one to the other.
- ❑ Blue Coat recommends using `gzip` encoding (or allowing both `gzip` and `deflate`) when using the HTTP compression feature.
- ❑ If the ProxySG appliance receives unknown content encoding and if content transformation is configured (such as popup blocking), an error results.
- ❑ If the origin server provides compressed content with a different compression level than that specified in policy, the content is not re-compressed.
- ❑ If the ProxySG appliance compressed and cached content at a different compression level than the level specified in a later transaction, the content is not re-compressed.
- ❑ Parsing of container HTML pages occurs on the server side, so pipelining (prefetching) does not work when the server provides compressed content.

- ❑ Compressing a zip file breaks some browser versions, and compressing images does not provide added performance. For a current list of content types that are not compressed, refer to the Release Notes.
- ❑ All responses from the server can be compressed, but requests to the server, such as POST requests, cannot. Only 200 OK responses can be compressed.

Section H: Viewing HTTP/FTP Statistics

This section discusses the following topics:

- ❑ "HTTP/FTP History Statistics"
- ❑ "Viewing the Number of HTTP/HTTPS/FTP Objects Served"
- ❑ "Viewing the Number of HTTP/HTTPS/FTP Bytes Served" on page 203
- ❑ "Viewing Active Client Connections" on page 204
- ❑ "Viewing HTTP/HTTPS/FTP Client and Server Compression Gain Statistics" on page 204
- ❑ "Disabling the Proxy-Support Header" on page 207

HTTP/FTP History Statistics

The **HTTP/FTP History** tabs display bar graphs that illustrate the last 60 minutes, 24 hours, and 30 days for the number of objects served, bytes served, active clients, and client and server compression gain statistics associated with the HTTP, HTTPS, and FTP protocols. The overall client and server compression-gain statistics are displayed under **System Usage**.

Note: You can view current HTTP statistics through the CLI using the `show http-stats` command.


Viewing the Number of HTTP/HTTPS/FTP Objects Served

The **HTTP/HTTPS/FTP Objects** tab illustrates the device activity over the last 60 minutes, 24 hours, and 30 days. These charts illustrate the total number of objects served from either the cache or from the Web.

The maximum number of objects that can be stored on a ProxySG is affected by a number of factors, including the SGOS version it is running and the hardware platform series.

To view the number of HTTP/HTTPS/FTP objects served:

1. From the Management Console, select **Statistics > Protocol Details > HTTP/FTP History > HTTP/HTTPS/FTP Objects**.
2. Select the **Duration:** from the drop-down list.




3. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.

Viewing the Number of HTTP/HTTPS/FTP Bytes Served

The **HTTP/HTTPS/FTP Bytes** tab shows the sum total of the number of bytes served from the device over the last 60 minutes, 24 hours, and 30 days. The chart shows the total number of bytes for objects served by the device, including both cache hits and cache misses.

To view the number of HTTP/HTTPS/FTP bytes served:

1. From the Management Console, select **Statistics > Protocol Details > HTTP/FTP History > HTTP/HTTPS/FTP Bytes**.
2. Select the **Duration:** from the drop-down list.




3. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.

Viewing Active Client Connections

The HTTP/HTTPS/FTP Clients tab shows the maximum number of clients with requests processed over the last 60 minutes, 24 hours, and 30 days. This does not include idle client connections (connections that are open but that have not made a request). These charts allow you to monitor the maximum number of active clients accessing the ProxySG at any one time. In conjunction with the HTTP/HTTPS/FTP Objects and HTTP/HTTPS/FTP Bytes tabs, you can determine the number of clients supported based on load, or load requirements for your site based on a specific number of clients.

To view the number of active clients:

1. From the Management Console select **Statistics > Protocol Details > HTTP/FTP History > HTTP/HTTPS/FTP Clients**.
2. Select the **Duration:** from the drop-down list.



3. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.

Viewing HTTP/HTTPS/FTP Client and Server Compression Gain Statistics

Under HTTP/FTP History, you can view HTTP/FTP client and server compression-gain statistics for the ProxySG over the last 60 minutes, 24 hours, and 30 days in the Client Comp. Gain and the Server Comp. Gain tabs. Overall client and server compression-gain statistics are displayed under System Usage. These statistics are not available through the CLI.

The green display on the bar graph represents uncompressed data; the blue display represents compressed data. Hover your cursor over the graph to see the compressed gain data.

See one of the following sections for more information:


- ["Viewing HTTP/FTP Client Compressed Gain Statistics"](#)

- ❑ ["Viewing HTTP/FTP Server Compressed Gain Statistics"](#) on page 205

Viewing HTTP/FTP Client Compressed Gain Statistics

To view HTTP/FTP client compressed gain statistics:

1. From the Management Console, select **Statistics > Protocol Details > HTTP/FTP History > Client Comp. Gain**.
2. Select the **Duration:** from the drop-down list.



3. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.


See Also

["Viewing HTTP/HTTPS/FTP Client and Server Compression Gain Statistics"](#) on page 204

Viewing HTTP/FTP Server Compressed Gain Statistics

To view HTTP/FTP server compressed gain statistics:

1. From the Management Console, select **Statistics > Protocol Details > HTTP/FTP History > Server Comp. Gain**.
2. Select the **Duration:** from the drop-down list.



- (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.

See Also

"Viewing HTTP/HTTPS/FTP Client and Server Compression Gain Statistics" on page 204

Section I: Supporting IWA Authentication in an Explicit HTTP Proxy

Internet Explorer does not allow IWA authentication through a ProxySG when explicitly proxied. To facilitate this authentication, Blue Coat added a `Proxy-Support: Session-based-authentication` header. By default, when the ProxySG receives a 401 authentication challenge from upstream, it sends the `Proxy-Support: Session-based-authentication` header in response.

The `Proxy-Support` header is not supported if:

- ❑ you are using an older browser (Refer to the *SGOS Release Notes* for supported browser versions).
- ❑ both the ProxySG and the OCS perform IWA authentication.

In either case, Symantec recommends that you disable the header and enable **Force IWA for Server Authentication**. The **Force IWA for Server Authentication** action converts the 401-type server authentication challenge to a 407-type proxy authentication challenge that Internet Explorer supports. The ProxySG also converts the resulting Proxy-Authentication headers in client requests to standard server authorization headers, which allows an IWA authentication challenge to pass through when Internet Explorer is explicitly proxied through the appliance.

Disabling the Proxy-Support Header


The Proxy-Support header is sent by default when an explicitly configured ProxySG receives a 401 authentication challenge from upstream.

The header modification policy allows you to suppress or modify the Proxy-Support custom header, and prevents the ProxySG from sending this default header. Use either the Visual Policy Manager (VPM) or CPL to disable the header through policy. For complete information on using VPM, refer to *Visual Policy Manager Reference*.

Note: To suppress the `Proxy-Support` header globally, use the `http force-ntlm` command to change the option. To suppress the header only in certain situations, continue with the procedures below.

To suppress the proxy-support header through the VPM:

1. In a Web Access Layer, right click in the **Action** field and select **Set**. The Set Action dialog displays.
2. Click **New** to see the drop-down list; select **Control Response Header**.



3. Fill in the fields as follows:
 - a. **Name:** Enter a meaningful name.
 - b. **Show:** Select **Custom** from the drop-down list.
 - c. **Header Name:** Enter Proxy-Support.
 - d. Verify **Suppress** is selected.
4. Click **OK**.
5. Click **Apply**.

To suppress the proxy-support header through CPL:

Use CPL to define the Proxy-Support custom header object and to specify what action to take. The example below uses Proxy-Support as the action name, but you can choose any name meaningful to you. The result of this action is to suppress the Proxy-Support header.

```
<Proxy>
  action.Proxy-Support (yes)

define action Proxy-Support
  delete(response.x_header.Proxy-Support)
end action Proxy-Support
```

Section J: Supporting Authentication on an Upstream Explicit Proxy

Proxy chaining may cause issues in HTTPS configurations. When an upstream proxy requires Proxy-Authentication, a timeout may occur because by the time the proxy authentication challenge occurs in the HTTP CONNECT request, the client has already established a non-authorized connection to the downstream proxy (which may or may not be a ProxySG).

Deployment Scenarios

Use this configuration when the ProxySG is inserted between a client and an explicit proxy configured to use authentication. It can also be helpful in transparent deployments.

- Explicit downstream: The ProxySG supports authentication to the client for SSL/HTTPS traffic, with an upstream proxy performing the authentication. The upstream proxy is not in your (control)
- Transparent downstream: The ProxySG, deployed transparently, supports authentication to the client for SSL/HTTPS traffic, with an upstream proxy performing the authentication. For example, in a chain where two proxies are configured transparently as accelerators and a third further upstream functions explicitly, authentication requests may not reach their destinations.

Section K: Detect and Handle WebSocket Traffic

(Introduced in SGOS 6.5.5.7) The Internet Engineering Task Force (IETF) standardized the WebSocket protocol in 2011. WebSocket provides simultaneous two-way communications channels over a single TCP connection by detecting the presence of a proxy server and tunneling communications through the proxy.

To upgrade an HTTP connection to a newer HTTP version or use another protocol such as WebSocket, a client sends a request with `Upgrade`, `Connection`, and other relevant headers. Previous versions of SGOS did not allow WebSocket handshakes to complete, but supported versions allow the handshake to complete successfully. This version also detects WebSocket traffic and allows you to perform specific policy actions.

When the appliance detects a WebSocket request in the HTTP/S request, the Active Sessions tab in the Management Console indicates that the traffic is WebSocket. Use the filter **Protocol > WebSocket**.

To differentiate WebSocket traffic in the access-log, use the `TCP_WEBSOCKET` value in the `s-action` field. You can determine if the traffic was plain WebSocket or secure WebSocket by looking at the scheme (HTTP or HTTPS).

How the ProxySG Appliance Handles an Upgrade Request

Refer to the following overviews of how the ProxySG appliance handles a WebSocket upgrade request in transparent proxy and in explicit proxy. For more information on the policy condition mentioned in the following overviews, refer to the *Content Policy Language Reference* and the *Visual Policy Manager Reference*.

Upgrade Request in Transparent Mode

- a. The browser sends a protocol upgrade request to the proxy.
- b. The HTTP proxy receives the upgrade request.
- c. If the `Upgrade` header has a single value of `websocket`, the HTTP proxy begins a WebSocket handshake by forwarding the `Upgrade` and `Connection` headers upstream to upgrade the connection protocol.

In this case, the `tunneled=yes` and `http.websocket=yes` conditions evaluate to true.

- d. Policy runs and evaluates the request. If the request is allowed, the proxy takes the next step depending on the response code:
 - If the HTTP response code is 101 ("Switching Protocols"), the proxy tunnels the request.
 - If the HTTP response code is successful (2xx), the proxy returns a **400 Bad Request** exception to indicate that the origin content server (OCS) did not understand the upgrade request.
 - In all other cases, the proxy returns the standard HTTP response codes and does not tunnel the request.

Note: The appliance evaluates all policy that applies to a transaction during the initial upgrade request.

Upgrade Request in Explicit Mode

- a. The browser sends an HTTP CONNECT request to the proxy.
- b. The HTTP proxy receives the HTTP CONNECT request.
- c. If **Detect Protocol** is enabled on the HTTP proxy, the request is forwarded to the HTTP proxy.

If **Detect Protocol** is disabled on the HTTP proxy and policy does not allow HTTP CONNECT requests, the appliance treats the request as if `force_protocol(http)` were set in policy.

The request is thus forwarded to the HTTP proxy, allowing the appliance to evaluate policy on (and possibly allow) tunneled HTTP traffic, such as WebSocket requests, while blocking non-HTTP protocols sent over HTTP CONNECT.

If **Detect Protocol** is disabled on the HTTP proxy and HTTP CONNECT is allowed in policy, the request is TCP-tunneled.

- d. (If the protocol is secure WebSocket) If **Detect Protocol** for SSL is disabled, the request is TCP-tunneled. If **Detect Protocol** for SSL is enabled, the request is forwarded to the SSL proxy.

(On the SSL proxy) If HTTPS interception is disabled, the request is SSL-tunneled. If HTTPS interception is enabled, the request is forwarded to the HTTPS proxy.

The proxy detects the `Upgrade: websocket` header and begins a WebSocket handshake. The `tunneled=yes` and `http.websocket=yes` conditions evaluate to true. policy that applies to a transaction during the initial upgrade request.

Feature Limitations

- ❑ The appliance does not perform ICAP scanning (either REQMOD or RESPMOD) on transactions using the WebSocket protocol.
- ❑ You must import the appliance's signing certificate authority (CA) certificate into the browser to prevent a trust error from occurring when the appliance intercepts HTTPS and detects WebSocket over HTTPS.

Chapter 9: Managing the SSL Proxy

This chapter discusses the ProxySG SSL proxy.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ [Section A: "Intercepting HTTPS Traffic"](#) on page 217
- ❑ [Section B: "Configuring SSL Rules through Policy"](#) on page 227
- ❑ [Section C: "Viewing SSL Statistics"](#) on page 232
- ❑ [Section D: "Using STunnel"](#) on page 235
- ❑ [Section E: "Tapping Decrypted Data with Encrypted Tap"](#) on page 241
- ❑ [Section F: "Working with an HSM Appliance"](#) on page 244
- ❑ [Section G: "Advanced Topics"](#) on page 249

For information on Certificate Authority (CA) certificates, keyrings, and key pairs, see ["Authenticating a ProxySG"](#) on page 1289.

About the SSL Proxy

HTTPS traffic poses a major security risk to enterprises. Because the SSL content is encrypted, it cannot be monitored by normal means. This enables users to bring in viruses, access forbidden sites, or leak confidential business information over the HTTPS connection on port 443.

The SSL proxy intercepts, decrypts and re-encrypts HTTPS traffic (in explicit and transparent modes) so that security measures such as authentication, virus scanning, and URL filtering, and performance enhancements such as HTTP caching can be applied to HTTPS content. Additionally, the SSL proxy validates server certificates presented by various HTTPS sites at the gateway and offers information about the HTTPS traffic in the access log.

The SSL proxy tunnels all HTTPS traffic by default unless there is an exception, such as a certificate error or a policy denial. In such cases, the SSL proxy intercepts the SSL connection and sends an error page to the user. The SSL proxy also enables interception of HTTPS traffic for monitoring purposes.

The SSL proxy can perform the following operations while tunneling HTTPS traffic.

- ❑ Validate server certificates, including revocation checks using Certificate Revocation Lists (CRLs).
- ❑ Check various SSL parameters such as cipher and version.
- ❑ Log useful information about the HTTPS connection.

When the SSL proxy is used to intercept HTTPS traffic, it can also:

- ❑ Cache HTTPS content.
- ❑ Apply HTTP-based authentication mechanism.
- ❑ Send decrypted data to a configured ICAP Antivirus appliance for virus scanning and URL filtering.
- ❑ Apply granular policy (such as validating mime type and filename extension).

IPv6 Support

The SSL proxy is able to communicate using either IPv4 or IPv6, either explicitly or transparently.

In addition, for any service that uses the SSL proxy, you can create listeners that bypass or intercept connections for IPv6 sources or destinations.

Validating the Server Certificate

The SSL proxy can perform the following checks on server certificates:

- ❑ Verification of issuer signature.
- ❑ Verification of certificate dates.
- ❑ Comparison of host name in the URL and certificate (intercepted connections only).

Host names in server certificates are important because the SSL proxy can identify a Web site just by looking at the server certificate if the host name is in the certificate. Most content-filtering HTTPS sites follow the guideline of putting the name of the site as the common name in the server's certificate.

- ❑ Verification of revocation status.

To mimic the overrides supported by browsers, the SSL proxy can be configured to ignore failures for the verification of issuer signatures and certificate dates and comparison of the host name in the URL and the certificate.

The ProxySG trusts all root CA certificates that are trusted by Internet Explorer and Firefox. This list is updated to be in sync with the latest versions of IE and Firefox.

Checking CRLs

An additional check on the server certificate is done through Certificate Revocations Lists (CRLs). CRLs show which certificates are no longer valid; the CRLs are created and maintained by Certificate Signing Authorities that issued the original certificates.

Only CRLs that are issued by a trusted issuer can be used by the ProxySG. The CRL issuer certificate must exist as CA certificate on the ProxySG before the CRL can be imported.

The ProxySG allows:

- ❑ One local CRL per certificate issuing authority.
- ❑ An import of a CRL that is expired; a warning is displayed in the log.
- ❑ An import of a CRL that is effective in the future; a warning is displayed in the log.

Working with SSL Traffic

The STunnel (SSL interception and tunnel) configuration intercepts all SSL traffic, handing HTTPS traffic off to the HTTPS forward proxy for compression and acceleration. STunnel decrypted traffic may be tapped and read by a third party application such as Wireshark or Snort.

Recommendations for intercepting traffic include:

- ❑ Intercept non-HTTPS traffic for acceleration
- ❑ Intercept any SSL traffic for tap, when you don't know the application protocol over SSL
- ❑ The HTTPS information in the next section applies as well.

Determining What HTTPS Traffic to Intercept

The SSL proxy tunnels HTTPS traffic by default; it does not intercept HTTPS traffic.

Many existing policy conditions, such as destination IP address and port number, can be used to decide which HTTPS connections to intercept.

Additionally, the SSL proxy allows the host name in the server certificate to be used to make the decision to intercept or tunnel the traffic. The server certificate host name can be used as is to make intercept decisions for individual sites, or it can be categorized using any of the various URL databases supported by Blue Coat.

Categorization of server certificate host names can help place the intercept decision for various sites into a single policy rule.

Recommendations for intercepting traffic include:

- ❑ Intercept Intranet traffic.
- ❑ Intercept suspicious Internet sites, particularly those that are categorized as *none* in the server certificate.

Recommendations for traffic to not intercept includes sensitive information, such as personal financial information.

Managing Decrypted Traffic

After the HTTPS connection is intercepted, you can do:

- ❑ Anti-virus scanning over ICAP.

- ❑ URL filtering (on box and off-box). Symantec recommends on box URL/content filtering if you use transparent proxy. When the URL is sent off-box for filtering, only the host name or IP address of the URL (not the full path) is sent for security reasons.
- ❑ Filtering based on the server certificate host name.
- ❑ Caching.

HTTPS applications that require browsers to present client certificates to secure Web servers do not work if you are intercepting traffic. To address this, you can create a policy rule to prevent the interception of such applications, or add client certificates to the ProxySG appliance, and write policy to present the correct certificate.

If you configure the ProxySG to intercept HTTPS traffic, be aware that local privacy laws might require you to notify the user about interception or obtain consent prior to interception. You can option to use the HTML Notify User object to notify users after interception or you can use consent certificates to obtain consent before interception. The HTML Notify User is the easiest option; however, the ProxySG must decrypt the first request from the user before it can issue an HTML notification page.


Using the SSL Proxy with ADN Optimization

The SSL proxy itself can be used as a split proxy, which requires two SSL proxies, one at the branch and one at the core, working together. A *split proxy* can be configured (see below) to implement functionality that is not possible in a standalone proxy.

In this configuration, the SSL proxy supports ADN optimization on WAN networks, and SSL traffic performance can be increased through the byte caching capability offered. The branch proxy, which makes the decisions, is configured with both ADN optimization and SSL proxy functionality.

The *Concentrator* proxy (a ProxySG that provides access to data center resources) does not require any configuration related to the SSL proxy. It only requires the necessary ADN configuration for applying byte caching capabilities to intercepted SSL content.

No special configuration is required to the SSL proxy.



Section A: Intercepting HTTPS Traffic

Intercepting HTTPS traffic (by decrypting SSL connections at the ProxySG) allows you to apply security measures like virus scanning and URL filtering. See [“Configuring STunnel” on page 235](#) to intercept HTTPS using STunnel.

Configuration to intercept HTTPS traffic requires the following tasks:

- ❑ A ProxySG SSL license is required before you can make use of the SSL proxy for interception. This can be verified in the maintenance tab > licensing page.
- ❑ Determine whether you are using transparent or explicit mode. For information on explicit versus transparent proxies, see [“Explicit and Transparent Proxy” on page 99](#).
- ❑ Create an SSL service or HTTP/SOCKS services with protocol detection enabled, depending on whether you are using transparent or explicit mode. The Detect Protocol setting is disabled by default. For more information on creating an SSL service, skip to [“Configuring the SSL Proxy in Transparent Proxy Mode” on page 218](#).
- ❑ Create or import an issuer keyring, which is used to sign emulated server certificates to clients on the fly, allowing the SSL proxy to examine SSL content. For more information on creating an issuer keyring, see [“Specifying an Issuer Keyring and CCL Lists for SSL Interception” on page 220](#).
- ❑ (Optional) Use the **Notify User** object or client consent certificates to notify users that their requests are being intercepted and monitored. Whether this is required depends on local privacy laws. The ProxySG has to decrypt the first request from the user to issue an HTML notification page. If this is not desirable, use client consent certificates instead. For more information on configuring the **Notify User** policy, refer to the *Visual Policy Manager Reference*. For information on managing client consent certificates, see [“Using Client Consent Certificates” on page 221](#).
- ❑ Download CA certificates to desktops to avoid a security warning from the client browsers when the ProxySG is intercepting HTTPS traffic. For information, see [“Downloading an Issuer Certificate” on page 221](#).
- ❑ Using policy (VPM or CPL), create rules to intercept SSL traffic and to control validation of server certificates. By default, such traffic is tunneled and not intercepted. You must create suitable policy before intercepting SSL traffic. For more information on using policy to intercept SSL traffic, see [Section B: “Configuring SSL Rules through Policy” on page 227](#).
- ❑ Configure the Blue Coat AV or other third-party ICAP vendor, if you have not already done this. For more information on ICAP-based virus scanning, see [“Configuring Threat Protection” on page 445](#) (Blue Coat AV) and [“Malicious Content Scanning Services” on page 459](#).
- ❑ Configure the Blue Coat Web Filter or a third-party URL-filtering vendor, if you have not already done this. For more information on configuring WebFilter, see [“Filtering Web Content” on page 377](#).

- ❑ Configure Access Logging. For more information on configuring access logging, see ["Configuring Access Logging"](#) on page 617.
- ❑ Customize Exception Pages: To customize exception pages (in case of server certificate verification failure), refer to the Advanced Policy Tasks chapter, Section E, of the [Visual Policy Manager Reference](#).

Configuring the SSL Proxy in Transparent Proxy Mode


Proxy services are configured from the Management Console or the CLI. If using the SSL proxy in transparent mode, continue with this section.

If you are using the SSL proxy in explicit mode, you might need an HTTP proxy or a SOCKS proxy. For information on configuring an SSL proxy in explicit mode, see ["Configuring the SSL Proxy in Explicit Proxy Mode"](#) on page 220.

You can use a TCP Tunnel service in transparent mode to get the same functionality. A TCP tunnel service is useful when you have a combination of SSL and non-SSL traffic going over port 443 and you do not want to break the non-SSL traffic. The SSL service requires that all requests to its port be SSL.

To configure an SSL service in transparent proxy mode:

1. From the Management Console, select the **Configuration > Services > Proxy Services** tab.
2. Click **New**. The Edit Service dialog displays.



3. In the **Name** field, enter a meaningful name for this SSL proxy service.
4. From the **Service Group** drop-down list, select to which service this configuration applies. By default, **Other** is selected.
5. Select **SSL** from the **Proxy settings** drop-down list.
6. **TCP/IP Settings** option: The **Early Intercept** option cannot be changed for the SSL proxy service.
7. Select ADN options:
 - **Enable ADN.** Select this option to configure this service to use ADN. Enabling ADN does *not* guarantee the connections are accelerated by ADN. The actual enable decision is determined by ADN routing (for explicit deployment) or network setup (for transparent deployment).
 - The **Optimize Bandwidth** option is selected by default if you enabled WAN optimization during initial configuration. Clear the option if you are not configuring WAN optimization.
8. Create a new listener:
 - a. Click **New**; if you edit an existing listener, click **Edit**.
 - b. In the **Source address** area, the most common selection is **All**, which means the service applies to requests from any client (IPv4 or IPv6). You can, however, restrict this listener to a specific IPv4/IPv6 address or user subnet/prefix length.
 - c. Select a **Destination address** from the options. The correct selection might depend on network configuration. For overviews of the options, see ["About Proxy Services"](#) on page 110.
 - d. In the **Port Range** field, enter a single port number or a port range on which this application protocol broadcasts. For a port ranges, enter a dash between the start and end ports. For example: 8080–8085
 - e. In the **Action** area, select the default action for the service: **Bypass** tells the service to ignore any traffic matching this listener. **Intercept** configures the service to intercept and proxy the associated traffic.
 - f. Click **OK** to close the dialog. The new listener displays in the **Listeners** area.
9. Click **OK** to close the Edit Service dialog.
10. Click **Apply**.

Continue with ["Specifying an Issuer Keyring and CCL Lists for SSL Interception"](#) on page 220.

Configuring the SSL Proxy in Explicit Proxy Mode

The SSL proxy can be used in explicit mode in conjunction with the HTTP Proxy or SOCKS Proxy. You must create an HTTP Proxy service or a SOCKS Proxy service and use it as the explicit proxy from desktop browsers. You must also ensure that the detect-protocol attribute is enabled for these services.

When requests for HTTPS content are sent to either a SOCKS proxy or an HTTP proxy, the proxies can detect the use of the SSL protocol on such connections and enable SSL proxy functionality.

Continue with ["Specifying an Issuer Keyring and CCL Lists for SSL Interception"](#) on page 220.

Specifying an Issuer Keyring and CCL Lists for SSL Interception


The SSL proxy can emulate server certificates; that is, present a certificate that appears to come from the origin content server. In actuality, Blue Coat has emulated the certificate and signed it using the issuer keyring. By default only the `subjectName` and the expiration date from the server certificate are copied to the new certificate sent to the client. The ProxySG will emulate RSA certificates, matching the key size up to 2048 bits. DSA server certificates are emulated with 1024 bit RSA certificates.

Note: Only keyrings with both a certificate and a keypair can be used as issuer keyrings.

You can also change the CA Certificate Lists (CCLs) that contain the CAs to be trusted during client and server certificate validation. The defaults are adequate for the majority of situations. For more information about CCLs, see ["Authenticating a ProxySG"](#) on page 1289.

To specify the keyring and CCLs:

1. From the Management Console, select **Configuration > Proxy Settings > SSL Proxy**.



The screenshot shows the 'SSL Proxy' configuration page with the 'General Settings' tab selected. The settings are as follows:

Setting	Value
Issuer Keyring:	default
CCL for Client Certificates:	<All CA Certificates>
CCL for Server Certificates:	browser-trusted
Enable CRL on emulated certificates	☐
CRL distribution point host name:	10.9.46.164

2. **Issuer Keyring:** From the drop-down menu, select the keyring to use as the issuer keyring. Any keyring with both a certificate and a keypair in the drop-down menu can be used.
3. **CCL for Client Certificates:** Choose which CAs are trusted when the SSL proxy validates client certificates. The default is **<All CA Certificates>**.

4. **CCL for Server Certificates:** Choose which CAs are trusted when the SSL proxy validates server certificates. The CCL for server certificates is relevant even when SSL proxy is tunneling SSL traffic. The default is **browser-trusted**.
5. **Enable CRL on emulated certificates:** (Added in 6.5.9.10) Select this to enable a CRL distribution point field on emulated certificates. Particularly useful for Microsoft services.
6. **CRL distribution point host name:** (Added in 6.5.9.10) Enter the host name of the issuer CA.
7. Click **Apply**.

To configure policy, see ["Configuring SSL Rules through Policy"](#) on page 227.

Using Client Consent Certificates

The SSL proxy, in forward proxy deployments, can specify whether a client (typically a browser) certificate is required. These certificates are used for user consent, not for user authentication. Whether they are needed depends upon local privacy laws.

With client consent certificates, each user is issued a pair of certificates with the corresponding private keys. Both certificates have a meaningful user-readable string in the **common name** field. One certificate has a string that indicates grant of consent something like: "Yes, I agree to SSL interception". The other certificate has a common name indicating denial of consent, something like: "No, I do not agree to SSL interception".

Policy is installed on the ProxySG to look for these common names and to allow or deny actions. For example, when the string "Yes, I agree to SSL interception" is seen in the client certificate common name, the connection is allowed; otherwise, it is denied.

To configure client consent certificates:

1. Install the issuer of the client consent certificates as a CA certificate.
2. In VPM, configure the **Require Client Certificate** object in the **SSL Layer > Action** column.
3. Configure the **Client Certificate** object in the **Source** column to match common names.

Downloading an Issuer Certificate

When the SSL proxy intercepts an SSL connection, it presents an emulated server certificate to the client browser. The client browser issues a security pop-up to the end-user because the browser does not trust the issuer used by the ProxySG. This pop-up does not occur if the issuer certificate used by SSL proxy is imported as a trusted root in the client browser's certificate store.


The ProxySG makes all configured certificates available for download via its management console. You can ask end users to download the issuer certificate through Internet Explorer or Firefox and install it as a trusted CA in their browser of choice. This eliminates the certificate popup for emulated certificates.

To download the certificate through Internet Explorer, see "[To download a certificate through Internet Explorer](#):" on page 222. To download a certificate through Firefox, see "[To download a certificate through Firefox](#):" on page 223.


To download a certificate through Internet Explorer:

Note: You can e-mail the console URL corresponding to the issuer certificate to end users so that the he or she can install the issuer certificate as a trusted CA.

1. Select the **Statistics > Advanced** tab.
2. Select **SSL**.
3. Click **Download a Certificate as a CA Certificate**; the list of certificates on the system display.
4. Click a certificate (it need not be associated with a keyring); the File Download Security Warning displays asking what you want to do with the file.



5. Click **Save**. When the **Save As** dialog displays, click **Save**; the file downloads.
6. Click **Open** to view the Certificate properties; the Certificate window displays.



7. Click the **Install Certificate** button to launch the **Certificate Import Wizard**.
8. Ensure the **Automatically select the certificate store based on the type of certificate** radio button is enabled before completing the wizard
9. Click **Finish**. the wizard announces when the certificate is imported.
10. (Optional) To view the installed certificate, go to Internet Explorer, **Select Tools > Internet Options > Contents > Certificates**, and open either the **Intermediate Certification Authorities** tab or the **Trusted Root Certification Authorities** tab, depending on the certificate you downloaded.

To download a certificate through Firefox:

Note: You can e-mail the console URL corresponding to the issuer certificate to end users so that the end-user can install the issuer certificate as a trusted CA.

1. Select the **Statistics > Advanced** tab.
2. Select **SSL**.
3. Click **Download a ProxySG Certificate as a CA Certificate**; the list of certificates on the system display.
4. Click a certificate (it need not be associated with a keyring); the **Download Certificate** dialog displays.



5. Enable the options needed. View the certificate before trusting it for any purpose.
6. Click **OK**; close the Advanced Statistics dialog.

Warn Users When Accessing Websites with Untrusted Certificates

Preserve Untrusted Certificate Issuer allows the ProxySG appliance to present the browser with a certificate that is signed by its untrusted issuer keyring. The browser displays certificate information to the user, and lets the user accept the security risk of an untrusted certificate and proceed to the website.

The `default-untrusted` keyring has been added to the ProxySG appliance to use with the Preserve Untrusted Certificate Issuer feature. The `default-untrusted` keyring should not be added to any trusted CA lists.

Note: This only applies to SSL forward proxy transactions with HTTPS interception enabled.

To display a warning to users about untrusted certificates on website, you must complete the following tasks.

Task #	Reference
1	"Presenting Untrusted Certificates to a Browser" on page 225
2	"Set the Behavior when Encountering Untrusted Certificates" on page 225

Presenting Untrusted Certificates to a Browser

Configure the ProxySG appliance to act as a certificate authority and present a certificate signed by a specific keyring for all traffic. The default is the `default-untrusted` keyring.

1. From the Management Console, select **Configuration > Proxy Settings > SSL Proxy**.
2. To have the ProxySG appliance act as a Certificate Authority (CA) and present the browser with an untrusted certificate, select **Preserve untrusted certificate issuer**.
3. From the **Untrusted Issuer Keyring** drop-down, select the desired keyring from the list of eligible keyrings which will be used to sign untrusted server certificates presented by the ProxySG appliance.
4. Click **Apply**.

Set the Behavior when Encountering Untrusted Certificates

In the VPM or CPL, define what the ProxySG appliance should do for specific traffic if the user tries to access a website with an untrusted certificate.

Define Behavior in the Visual Policy Manager (VPM)

Override the ProxySG Management Console settings for specific traffic, to specify whether the users should be prompted when a certificate that has not been signed by a trusted Certificate Authority is encountered.

In the SSL Intercept Layer, add one of the following Actions:

❑ **Do not Preserve Untrusted Issuer**

If an OCS presents a certificate to the ProxySG appliance that is not signed by a trusted Certificate Authority (CA), the ProxySG appliance either sends an error message to the browser, or ignores the error and processes the request, based on the configuration of the Server Certificate Validation object.

❑ **Preserve Untrusted Issuer**

If an OCS presents a certificate to the ProxySG appliance that is not signed by a trusted Certificate Authority (CA), the ProxySG appliance acts as a CA and presents the browser with an untrusted certificate. A warning message is displayed to the user, and they can decide to ignore the warning and visit the website or cancel the request.

❑ **Use Default Setting for Preserve Untrusted Issuer**

The **Preserve untrusted certificate issuer** configuration setting in the ProxySG Management Console is used to determine whether or not untrusted certificate issuer should be preserved for a connection. This is the default behavior.

Define Behavior in CPL

Include the following syntax in policy to specify the behavior of the ProxySG appliance when users encounter a website with an untrusted certificate:

```
ssl.forward_proxy.preserve_untrusted(auto|yes|no)
```

where:

- **auto** - Uses the **Preserve untrusted certificate issuer** configuration setting in the ProxySG Management Console to determine whether untrusted certificate issuer should be preserved for a connection. This is the default.
- **yes** - Preserve untrusted certificate issuer is enabled for the connection.
- **no** - Preserve untrusted certificate issuer is disabled for the connection.

For example, to use the enable using the preserve untrusted certificate issuer, use the following syntax:

```
<ssl-intercept>  
  ssl.forward_proxy.preserve_untrusted(yes)
```


Section B: Configuring SSL Rules through Policy

SSL interception and access rules, including server certificate validation, are configured through policy—either the VPM or CPL. Use the **SSL Intercept Layer** to configure SSL interception; use the **SSL Access Layer** to control other aspects of SSL communication such as server certificate validation and SSL versions. To configure SSL rules using CPL, refer to the *Content Policy Language Reference*. This section covers the following topics:

- ❑ ["Using the SSL Intercept Layer" on page 227.](#)
- ❑ ["Using the SSL Access Layer" on page 229](#)
- ❑ ["Using Client Consent Certificates" on page 221](#)

The policy examples in this section are for in-path deployments of ProxySG appliances.

Using the SSL Intercept Layer

The SSL intercept layer allows you to set intercept options:

- ❑ ["To intercept HTTPS content through VPM:" on page 227](#)
- ❑ ["To intercept HTTPS requests to specific sites through the VPM:" on page 228](#)
- ❑ ["To customize server certificate validation through VPM:" on page 229](#)
- ❑ ["Configuring STunnel" on page 235](#)

Note: For detailed instructions on using VPM, refer to the *Visual Policy Manager Reference*.

To intercept HTTPS content through VPM:

1. Select the **Configuration > Policy > Visual Policy Manager** tab and launch the VPM.
2. From the **Policy** drop-down menu, select **Add SSL Intercept Layer**.
3. Right-click **Set** in the **Action** column; the **Set Action** object displays.
4. Click **New** and select **Enable HTTPS Intercept** object.

The options for **Issuer Keyring**, **Hostname**, **Splash Text**, and **Splash URL** all control various aspects for certificate emulation. Fill in the fields as follows:

- a. **Issuer Keyring:** If you selected an issuer keyring previously, that keyring displays. If you did not select an issuer keyring previously, the default keyring displays. To change the keyring that is used as the issuer keyring, choose a different keyring from the drop-down menu.
- b. **Hostname:** The host name you put here is the host name in the emulated certificate.
- c. **Splash Text:** You are limited to a maximum of 200 characters. The splash text is added to the emulated certificate as a certificate extension.

- d. **Splash URL:** The splash URL is added to the emulated certificate as a certificate extension.

The STunnel options control various aspects of SSL interception.


- a. **Enable STunnel Interception:** Establish a policy where configured STunnel services (such as POP3S and SMTPS) are terminated and accelerated.
- b. **Enable SSL interception with automatic protocol detection:** In addition to STunnel interception as described above, discovered HTTPS is handed off to the HTTPS proxy. Otherwise, SSL traffic continues in STunnel mode.

5. Click **OK** to save the changes.

You can use the **Disable SSL Intercept** object to disable HTTPS Intercept.

To intercept HTTPS requests to specific sites through the VPM:


1. Select the **Configuration > Policy > Visual Policy Manager** tab and launch the VPM.
2. From the **Policy** drop-down menu, select **Add SSL Intercept Layer**.
3. In the **Destination** column, right-click **Set**; the **Set Destination Object** displays.
4. Click **New** and select **Server Certificate**.



5. Fill in the fields as described below. You can only select one field:
 - a. **Hostname:** This is the host name of the server whose traffic you want to intercept. After entering the host name, use the drop-down menu to specify **Exact Match**, **Contains**, **At Beginning**, **At End**, **Domain**, or **Regex**.
 - b. **Subject:** This is the subject field in the server's certificate. After you enter the subject, use the drop-down menu to specify **Exact Match**, **Contains**, **At Beginning**, **At End**, **Domain**, or **Regex**.
6. Click **Add**, then **Close**; click **OK** to add the object to the rule.

To categorize host names in server certificates through VPM:

1. While still in the **Destination** column of the **SSL Intercept** layer, right-click **Set**; the **Set Destination** object displays.
2. Click **New** and select the **Server Certificate Category** object. The **Add Server Certificate Category Object** displays. You can change the name in the top field if needed.



3. Select the categories. The categories you selected display in the right-hand column.
4. Click **OK**.

Using the SSL Access Layer


For a list of the conditions, properties, and actions that can be used in the **SSL Access Layer**, refer to the *Content Policy Language Reference*.

Note: For detailed instructions on using VPM, refer to the *Visual Policy Manager Reference*.

To customize server certificate validation through VPM:

Note: The policy property `server.certificate.validate`, if set, overrides the `ssl-verify-server` command for either HTTP or for forwarding hosts.

1. Select the **Configuration > Policy > Visual Policy Manager** tab and launch the VPM.
2. From the **Policy** drop-down menu, select **Add SSL Access Layer**.
3. In the **Action** column, right-click **Set**; the **Set Action** object displays.
4. Click **New** and select **Set Server Certificate Validation** object.



5. By default, server certificate validation is enabled; to disable it, select **Disable server certificate validation** at the bottom of the dialog.

If server certificate validation is enabled, you can determine behavior by selecting the **Ignore hostname mismatch**, **Ignore expiration**, or **Ignore untrusted issuer** options. These options mimic the overrides supported by most browsers.

6. Select an option for revocation checks:
- Select an Online Certificate Status Protocol (OCSP) option. For more information, see [Section F: "Checking Certificate Revocation Status in Real Time \(OCSP\)"](#) on page 1151.
 - **Use only local certificate revocation check:** Uses the CRL configured on the ProxySG to perform the revocation check for a server certificate.
 - **Do not check certificate revocation:** Does not check the revocation status of the server certificate; however it still carries out the other certificate validation checks.
7. Click **OK**; click **OK** again to add the object.

Notes

Note: Pipelining configuration for HTTP is ignored for HTTPS requests intercepted by the SSL proxy. When the SSL proxy intercepts an HTTPS request, and the response is an HTML page with embedded images, the embedded images are not pre-fetched by the ProxySG.

- ❑ If the ProxySG and the origin content server cannot agree on a common cipher suite for intercepted connections, the connection is aborted.
- ❑ Server-Gated Cryptography and step-up certificates are treated just as regular certificates; special extensions present in these certificates are not be copied into the emulated certificate. Clients relying on SGC/step-up certificates continue using weaker ciphers between the client and the ProxySG when the SSL proxy intercepts the traffic.

Section C: Viewing SSL Statistics

The following sections discuss how to analyze various statistics generated by SSL transactions.

Viewing SSL History Statistics

The **Statistics > Protocol details > SSL History** tabs (**Unintercepted SSL Data**, **Unintercepted SSL Clients**, **Unintercepted SSL Bytes**) provide various useful statistics for unintercepted SSL traffic.

Note: Some SSL statistics (SSL client connections and total bytes sent and received over a period of time) can only be viewed through the Management Console (see "[Unintercepted SSL Data](#)" on page 232 and "[Unintercepted SSL Clients](#)" on page 233).

Unintercepted SSL Data

The **Unintercepted SSL Data** tab on the Management Console displays SSL statistics. The following table details the statistics provided through the **Unintercepted SSL Data** tab.

Table 9–1 Unintercepted SSL Data Statistics

Status	Description
<i>Current Unintercepted SSL Sessions</i>	The current number of unintercepted SSL client connections.
<i>Total Unintercepted SSL Sessions</i>	The cumulative number of unintercepted SSL client connections since the ProxySG was last rebooted.
<i>Total Bytes Sent</i>	The total number of unintercepted bytes sent.
<i>Total Bytes Received</i>	The total number of unintercepted bytes received.

To view unintercepted SSL data statistics:

From the Management Console, select the **Statistics > Protocol Details > SSL History > Unintercepted SSL Data** tab.

The default view shows all unintercepted SSL data.




Unintercepted SSL Clients

The **Unintercepted SSL Clients** tab displays dynamic graphical statistics for connections received in the last 60-minute, 24-hour, or 30-day period.

To view SSL client unintercepted statistics:

1. From the Management Console, select the **Statistics > Protocol Details > SSL History > Unintercepted SSL Clients** tab.




2. Select a time period for the graph from the **Duration:** drop-down list. The default is **Last Week**.
3. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.

Unintercepted SSL Bytes

The **Unintercepted SSL Bytes** tab displays dynamic graphical statistics for bytes received in the last 60-minute, 24-hour, or 30-day period.

To view unintercepted SSL byte statistics:

1. From the Management Console, select the **Statistics > Protocol Details > SSL History > Unintercepted SSL Bytes** tab.



2. Select the **Duration:** for the graph from the drop-down list. The default is Last week.
3. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.

Section D: Using STunnel

STunnel intercepts SSL traffic regardless of the application protocol over it. HTTPS traffic may be identified and handed off to that proxy, and you may create services to inspect and accelerate other SSL protocols, such as SMTPS. The decrypted data may be tapped; see ["Tapping Decrypted Data with Encrypted Tap"](#) on page 241.

STunnel integrates with secure ADN. When secure ADN is enabled, SSL traffic is accelerated using byte-caching and/or compression. An STunnel service will intercept traffic based on the configuration and policy. For intercepted SSL-sessions, the STunnel proxy acts as man-in-the-middle.

The STunnel sub-proxy can perform the following actions:

- ❑ Intercept SSL traffic and hands off HTTPS content to the HTTPS proxy when it is detected.
- ❑ Intercept non-HTTPS traffic.
- ❑ With ADN, accelerate intercepted SSL traffic.

If you are familiar with configuring an inline or explicit HTTPS proxy, STunnel works the same way. STunnel is configured with the policy rule

```
ssl.forward_proxy(yes) or ssl.forward_proxy(stunnel).
```

Traffic is handled by STunnel, and tunneled through or processed as appropriate.

STunnel supports SSLv2, SSLv3, TLS 1.0, TLS 1.1 and TLS 1.2.

Configuring STunnel


You can configure STunnel using the Visual Policy Manager (VPM) or the Management Console.

Configure STunnel Policy using VPM

STunnel, which lets you intercept SSL traffic regardless of the application protocol over it, is configured on the interception layer. To configure STunnel policy using the VPM, follow these steps.

1. On the **Configuration** tab, select **Policy > Visual Policy Manager**, then click **Launch**.
2. Select **Policy > Add SSL Intercept Layer**.
3. Right click in the **Action** column
4. Choose **Set > New > Enable SSL Interception**.
5. On the **Add SSL Interception Object** window, choose one of the following:
 - **Enable STunnel Interception:** Establish a policy where configured STunnel services (such as POP3S and SMTPS) are terminated and accelerated. Make sure to configure the related services if you choose this option.

- **Enable SSL interception with automatic protocol detection:** In addition to STunnel interception as described above, discovered HTTPS is handed off to the HTTPS proxy. Otherwise, SSL traffic continues in STunnel mode.



6. Click **OK**. The window closes.
7. Click **OK** on the **Set Action Object Window**; it closes.
8. To examine the policy, press **View**.
9. Click **Install policy** on the VPM window.

Configure STunnel via the Management Console

Accelerate SSL Traffic

To provide acceleration to SSL using byte caching, enable a secure ADN. See ["Using the SSL Proxy with ADN Optimization"](#) for details.

Intercept the traffic as described in ["Intercept SSL Based Traffic"](#).

For an inline or explicit forward proxy, use the policy rule `ssl.forward_proxy(stunnel)` or `ssl.forwrd_proxy(yes)`.

Note: If an unsuccessful SSL interception occurs (the SSL handshake fails), the traffic is tunneled.

Intercept SSL Based Traffic

Use STunnel to intercept SSL traffic, such as POP3S, SMTPS, and HTTPS.

Configure a Forward Proxy with an STunnel Service

1. Setup a Secure ADN between the concentrator and branch peer. See ["Verify Secure ADN"](#) on page 284.

2. On the branch peer, edit or create PO3S or SMTPS services (create a new service at **Configuration > Services > Proxy Services > New Service**).
3. Click **Apply** on the **Configuration** tab.

Name: POP3S-2

Service Group: Standard

Proxy settings

Proxy: SSL

☒ Detect Protocol

TCP/IP Settings

☒ Early Intercept

Application Delivery Network Settings

☒ Enable ADN

☒ Enable byte caching

Retention priority: normal

☒ Enable compression

Source IP	Destination IP	Port range	Action
All	All	995	Intercept

New Edit Delete

OK Cancel

Example POP3S Setup:

POP3S is located in the Standard group by default.

Name: POP3S

Service Group: Standard

Proxy Settings/Proxy: SSL

Detect Protocol: Check; identified HTTPS traffic will be handed to the HTTPS forward proxy for processing

TCP/IP: N/A

Application Delivery Network Settings: Click **Enable ADN**; **Retention** priority is set to **normal**.

Listeners: Set **Action** to **Intercept**.

For an SMTPS setup, follow the same configuration, save for choose the appropriate port and enter **SMTPS** as the **Name**.

Make sure your SSL policy is configured correctly for STunnel. See the next section.

Viewing STunnel Results

Traffic and Service results are available for STunnel. See ["Statistics"](#) on page 667 for additional details on understanding the presentation of statistics.


- STunnel is part of the SSL proxy, but is broken out under STunnel in the Proxy statistics, so you may easily find the results.
- The SSL proxy controls the tunneled (unintercepted) SSL traffic—there is no need to look at the Bandwidth Savings report for the SSL proxy since this traffic is not accelerated.
- For a typical setup, where you have HTTPS traffic identified and handed off to its proxy, make sure to look at the HTTPS proxy statistics and bandwidth savings as well as STunnel reports in order to get the best understanding of STunnel results.

Viewing Traffic Statistics

To see traffic statistics, go to **Statistics > Traffic Details**. View the **Traffic Mix** and **Traffic History** tab statistics. STunnel sessions are listed under **Active** and **Errored** sessions, as well.

Traffic Mix


On the **Traffic Mix > Service** tab, view traffic distribution and bandwidth statistics for SSL service traffic running through the ProxySG.



Traffic History

STunnel sessions are listed under **Traffic Mix** and **Traffic History**:

1. Select the **Statistics > Traffic Details > Traffic Mix/Traffic History**.
2. On the **Traffic Mix** tab, select **Proxy**.
 - The **BW Usage** and **BW Gain** tabs are available.
 - The pie chart visually represents the bandwidth percentage for each proxy, including STunnel.
 - Scroll down in the table to view the STunnel (and HTTPS) information.
3. On the **Traffic History** tab, select **Proxy**.
 - View **STunnel** on the **BW Usage**, **BW Gain**, **Client Bytes** and **Server Bytes** tabs.




Application Mix

(Starting in SGOS 6.5.6.1) The ProxySG appliance can classify SSL-tunneled traffic without full HTTPS interception. The **Statistics > Application Details > Application Mix** and **Statistics > Application Details > Application History** reports display the applications detected in SSL-tunneled traffic. In the Proxy Type column in Application Mix report, look for **STunnel**.

Viewing Session Statistics

To see STunnel accelerated session statistics such as duration, bandwidth savings using ADN functionality, and caching for current active and historical errored sessions, view the **Sessions** statistics on the **Statistics** tab.

1. On the Concentrator peer, log in to the Management Console.
2. Select the **Statistics > Sessions > Active Sessions/Errored Sessions > Proxied Sessions** tab.
3. From the **Filter** drop-down list, select **Proxy**.
4. Select **STunnel** from the corresponding drop down list.
5. Press **Show**.




See "Active Sessions—Viewing Per-Connection Statistics" on page 690 for details on using these windows.

Viewing Protocol Details

Go to **Protocol Details > SSL Data** tab to view client connection and data transfer bytes information for STunnel.

At **Protocol**, select **STunnel**.



Access Logging

View the SSL log to see the STunnel sessions; the **cs-protocol** value is set to **stunnel**.

Section E: Tapping Decrypted Data with Encrypted Tap

Encrypted tap streams decrypted data from intercepted HTTPS or STunnel SSL transactions on client connections. The tap is performed simultaneously and on the same ProxySG appliance which is performing the Secure Web Gateway function. The data is presented in a format that can be understood by common network traffic analysis tools like Wireshark, common network intrusion detection systems such as Snort, and so on.

- Encrypted Tap does not support VLAN.
- MTU is fixed at 1500 bytes.
- SSL protocol headers/records/details are not preserved.
- Encrypted Tap is supported for forward proxy for STunnel and HTTPS, and for reverse proxy for HTTPS.
- (Introduced in SGOS 6.5.5.7) Encrypted tap also taps WebSocket.

Before you start


- Ensure your SGOS license is up to date and includes a valid Encrypted Tap component
- Configure HTTPS (see [“Intercepting HTTPS Traffic” on page 217](#)) or STunnel (see [“Using STunnel” on page 235](#)) interception on the ProxySG appliance.
- Ensure the ProxySG appliance has at least one open Ethernet port.
- Have a computer with a spare, unused/assigned Ethernet interface and a third party analysis application installed available to receive the tapped data.

Follow these steps

On the ProxySG appliance:

1. Enable Proxy Services for HTTPS/HTTP:
 - a. From the Management Console, select **Configuration tab > Services > Proxy-Services**.
 - b. On the **Proxy Services** tab, select **Predefined Service Group >Standard > HTTPS**, and press **Edit Service**.
 - c. On the **Edit Service** pop up, under **Listeners**, set **Action=Intercept**.
 - d. Click **OK**. From the Management Console, select **Configuration > Services > Proxy Services**. **Edit Service** pop up closes.
2. On the **Configuration tab > Proxy Settings > General > General** tab, check **Reflect Client IP** to reflect the client IP.
3. From the Management Console, select **Configuration tab > Policy > Policy Options > Default Proxy Policy: Allow** to set the Default Policy to Allow.
4. Create the Encrypted Tap policy.

- a. From the Management Console, on the **Configuration** tab, select **Policy > Visual Policy Manager > Launch**. The **Visual Policy Manager** window pops up.
- b. On the **VPM**, from **Policy**, select **Add SSL Access Layer**, and provide a name as required.
- c. Highlight the added row, right click on **Action**, and choose **Set**.
- d. On the **Set Action Object** window, click **New...**, and choose **Enable encrypted tap**.
- e. On the **Add Encrypted Tap Object** window, set the name, verify **Enable encrypted tap** is selected, and choose the tap **Interface** to use from the drop down.



- f. Press **Ok**. The window closes.
 - g. Press **Ok**. The **Set Action Object** window closes.
 - h. Press **Install Policy**. You will see a confirmation when the new policy has been installed.
5. Install the Encrypted Tap policy.

Note: Make sure the tapped interface is not the same as any client/server/management interface in use, in order to avoid dumping tapped or decrypted traffic onto real servers. Furthermore, to avoid dropping traffic at the L2

device (resultant of how L2 forwarding works), ensure there are no Layer 2 bridging devices between the ProxySG appliance and the sniffer tools used on the tapped interface.

On another computer:

1. Connect the PC to the selected Ethernet interface.
2. Open the third-party application (such as Wireshark), and configure it to monitor the network traffic on the selected Ethernet interface. The intercepted HTTPS traffic should now be viewable by this application.

Viewing Encrypted Tap Results

❑ Tapping the Traffic

Traffic is accessed at the specified interface. It has a TCP-like format which networking monitoring tools such as Wireshark and Snort can easily interpret. Here are the output details:

- TCP-SYN/ACK for connection setup
 - TCP-FIN/ACK or TCP-RST for connection tear downs.
 - Original source and destination IP and ports of the connection
 - TCP sequence numbers, acknowledgements, and checksums, updated accordingly for data output
 - TTL set to 1
 - MAC addresses selected to avoid any potential conflicts. The Source MAC is the original source MAC address. If the Destination MAC address belongs to the original ProxySG, it may be translated, but will otherwise be preserved.
- ❑ View the **ssl log** to see HTTPS or STunnel sessions; tapped transactions have the `x-cs-connection-encrypted-tap` value set to **TAPPED**.


Troubleshooting

This section describes troubleshooting tips and solutions for Encrypted Tap.

- ❑ View access logs for Encrypted Tap, as shown above. See 'Viewing Access-Log Statistics.'
- ❑ View the Encrypted Tap debug log and statistics.
- ❑ Perform a packet capture at the hardware interface on the ProxySG. Go to **Maintenance > Service Information > Packet Captures** to access packet captures. The capture provides details on the data transmitted by the ProxySG; compare this to the received tap data.
- ❑ Perform policy tracing; refer to MySymantec for articles on how to perform an SSL policy trace.

Section F: Working with an HSM Appliance

A Hardware Security Module (HSM) provides additional security for storing cryptographic keys and certificates, which is required in some highly regulated industries. The Proxy SG Appliance can use a network-attached HSM appliance to store resigning CA keys, and to perform digital signature operations. The ProxySG Appliance exchanges signing requests and responses with the attached HSM appliance, over mutually authenticated HTTPS requests. The ProxySG Appliance sends certificate data to the HSM.



The ProxySG appliance can work with multiple HSM appliances, and multiple appliances can work with the same HSM. In the event that a policy rule using an HSM to sign cannot work due to lack of response from the HSM, the attempt is logged, and the appliance responds with an exception. In addition to the resigning certificates, a mutually authenticated connection (communication pipeline) must be set up by verified certificates.

Working with the SafeNet Java HSM

The SafeNet Java HSM must be configured separately. Additionally, Symantec provides an agent to install on the SafeNet Java HSM, which will be used to interact with other appliances (see the *BCHSM Agent Installation and Operations Guide*). A certificate to authorize the agent is included. The HSM Agent operates on top of a secure session. It communicates to the external Symantec entity (ProxySG appliance), and is used remotely.

Before You Begin

In order for the ProxySG to trust the HSM, you must import the server certificate for the HSM, and put it in to a CA Certificate List. Go to **Configuration > SSL > CA Certificates**, and **Import** the certificate. Name the certificate and paste the .PEM data in to the appropriate field. For further information, see ["Importing CA Certificates"](#).


An HSM requires a linked Device Profile (go to **SSL > Device Profiles**). Click **New**, and create a FIPS compliant or non-compliant profile as required, then enter the HSM credentials into the **Create SSL Device Profile** window. See ["Specifying an Issuer Keyring and CCL Lists for SSL Interception"](#) for more information.

Add an HSM

Click **HSM** in the **Configuration > SSL** menu. The new page contains three tabs.

Click **Create** on the **HSM** tab to set up an HSM connection.

1. On the **Configuration > SSL > HSM** tab, select **Create**. The **Create HSM** pane displays.
2. Enter the HSM credentials. For the **Device Profile**, select the HSM profile created earlier. Click **OK** to save the information and close the pane.
3. Click **Apply** on the **HSM** tab. The new HSM device appears in the list. **Referenced** will show “No” until you use the new HSM in policy.




Add an HSM Keyring

Adding an HSM keyring follows the same steps as adding any SSL keyring. HSM keyrings are also available in **Proxy Settings > SSL Proxy > General Settings > Issuer Keyring**.

1. On the **Configuration > SSL > HSM > HSM Keyrings** tab, select **Create**. The **Create HSM Keyring** window pops up.
2. Enter the HSM credentials. Select **Paste From Clipboard** button to copy in the **Certificate** .PEM file; the **Key Label** is the name associated with the private key created on the SafeNet Java HSM. Click **OK** to save the information and close the window.
3. Click **Apply** on the **HSM Keyrings** window. The new HSM keyring appears in the list. **Referenced** will show “No” until you use the new keyring in policy.

Note: A keyring which is referenced by policy can't be deleted.



Once a keyring has been created, you can click **View Certificate** to see the certificate details and .PEM file data. Click **Preview** to see a list of actions which will occur when the keyring is implemented.


Note: HSM keyrings also appear in the **Proxy Settings > SSL Proxy** list of **Issuer Keyrings**.

Adding an HSM Keygroup

Keygroups may be referenced in policy, instead of an individual keyring. When a keygroup is used, the SSL connections are load balanced, either within one HSM, or across an HSM group.

Adding an HSM keygroup follows the same steps as adding any SSL keylist (see "[Group Related Client Keyrings into a Keylist](#)").


1. On the **Configuration > SSL > HSM Keygroups** tab, select **Create**. The **Create HSM Keygroup** window pops up. Any preexisting keygroups appear in the **Available HSM Keyrings** fields.
2. Create the new group. Click the **Add>>** and **Remove>>** to move keyrings from the **Available HSM Keyrings** list to the **Included HSM Keyrings** list, to have them included in the new group.
3. Click **OK**. The window closes.
4. Click **Apply** on the **HSM Keygroups** window.



Write HSM Policy

Use policy to direct the SSL proxy to use an HSM keyring or keygroup to sign an emulated certificate from an intercepted authenticated SSL connection.

1. Launch the VPM (**Policy > Visual Policy Manager > Launch**).
2. On the **Blue Coat Visual Policy Manager** window, select **Policy > Add SSL Intercept Layer**.
3. Rename the layer on the **Add New Layer** window if required, then click **OK** (not shown in the graphic).
4. Highlight the new layer, and right click at **Action**; select **Set**. The **Set Action Object** window displays.
5. On the **Set Action Object** window, select **New. > Enable SSL Interception**.
6. On the **Add SSL Interception Object** window, select the **Issuer Keyring** to use for HSM signatures. Configured HSM keyrings and keygroups appear on the drop down list.
7. Click **OK**. The window closes.
8. Click **Install Policy**. You will see a "Policy installation was successful" message on completion.
9. Close the VPM and click **Apply**.



Section G: Advanced Topics

If you use OpenSSL or Active Directory, you can follow the procedures below to manage your certificates.

For OpenSSL, see "[Creating an Intermediate CA using OpenSSL](#)" on page 249; if using Active Directory, see "[Creating an Intermediate CA using Microsoft Server 2012 \(Active Directory\)](#)" on page 252.

Creating an Intermediate CA using OpenSSL

This section describes the certificate management when creating an intermediate CA using OpenSSL.

The overall steps are:

- ❑ "[Installing OpenSSL](#)" on page 249
- ❑ "[Creating a Root Certificate](#)" on page 249
- ❑ "[Modifying the OpenSSL.cnf File](#)" on page 250
- ❑ "[Signing the ProxySG CSR](#)" on page 250
- ❑ "[Importing the Certificate into the ProxySG](#)" on page 251
- ❑ "[Testing the Configuration](#)" on page 251

Various OpenSSL distributions can be found at <http://www.openssl.org>.

Installing OpenSSL

After OpenSSL is installed, you must edit the `openssl.cnf` file and ensure the path names are correct. By default root certificates are located under `./PEM/DemoCA`; generated certificates are located under `/certs`.

Creating a Root Certificate

In order to create a root Certificate Authority (CA) certificate, complete the following steps.

Note: The key and certificate in this example is located at `./bin/PEM/demoCA/private/`.

1. In command prompt, enter:

```
openssl req -new -x509 -keyout
c:\resources\ssl\openssl\bin\PEM\demoCA\private\
cakey.pem -out
c:\resources\ssl\openssl\bin\PEM\demoCA\private\CACert.pem
```

where the root directory for openssl is: `\resources\ssl\openssl`

```
openssl req -new -x509 -keyout
c:\resources\ssl\openssl\bin\PEM\demoCA\private\cakey.pem -out
c:\resources\ssl\openssl\bin\PEM\demoCA\private\CACert.pem
Using configuration from C:\Resources\SSL\OpenSSL\bin\openssl.cnf
Loading 'screen' into random state - done
```

```
Generating a 1024 bit RSA private key
.....+++++
.....+++++
writing new private key to
'c:\resources\ssl\openssl\bin\PEM\demoCA\private\cakey.pem'
Enter PEM pass phrase:
```

2. Type any string more than four characters for the PEM pass phrase.
3. Enter the certificate parameters, such as country name, common name that are required for a Certificate Signing Request (CSR).

The private key and root CA are now located under the directory `./PEM/DemoCA/private`

4. Create a ProxySG keyring.
 - a. From the Management Console, select **Configuration > SSL > Keyrings**.
 - b. Click **Create**; fill in the fields as appropriate.
 - c. Click **OK**.
5. Create a CSR on the ProxySG.
 - a. From the Management Console, select **Configuration > SSL > Keyrings**.
 - b. Highlight the keyring you just created; click **Edit/View**.
 - c. In the Certificate Signing Request pane, click **Create** and fill in the fields as appropriate.

Note: Detailed instructions on creating a keyring and a CSR are in ["Authenticating a ProxySG"](#) on page 1289. They can also be found in the online help.

6. Paste the contents of the CSR into a text file called `new.pem` located in the `./bin` directory.

Modifying the OpenSSL.cnf File

Modify the `openssl.cnf` file to import the openssl root CA into your browser. If you do not do this step, you must import the ProxySG certificate into the browser.

1. In the `openssl.cnf` file, look for the string `basicConstraints=CA`, and set it to `TRUE`.

```
basicConstraints=CA:TRUE
```
2. Save the `openssl.cnf` file.

Signing the ProxySG CSR

Open a MS-DOS window and enter:

```
openssl ca -policy policy_anything -out newcert.pem -in new.pem
```

The output is:


```
Using configuration from C:\Resources\SSL\OpenSSL\bin\openssl.cnf
Enter PEM pass phrase:
Check that the request matches the signature
Signature ok
The Subjects Distinguished Name is as follows
countryName             :PRINTABLE:'FR'
stateOrProvinceName     :PRINTABLE:'Paris'
localityName            :PRINTABLE:'Paris'
organizationName        :PRINTABLE:'BlueCoat'
organizationalUnitName  :PRINTABLE:'Security Team'
commonName              :PRINTABLE:'ProxySG.bluecoat.com'
emailAddress            :IA5STRING:'support@bc.com'
Certificate is to be certified until Sep 27 13:29:09 2006 GMT (365
days)
Sign the certificate? [y/n]:y
1 out of 1 certificate requests certified, commit? [y/n]y
Write out database with 1 new entries
Data Base Updated
```

This signs the certificate; it can then be imported into the ProxySG.

Importing the Certificate into the ProxySG

1. Open the file `newcert.pem` in a text editor.
2. Select **Management Console > Configuration > SSL > SSL Keyrings**.
3. Selecting the keyring used for SSL interception; click **Edit/View**.
4. Paste in the contents of the `newcert.pem` file.
5. Import the contents of the `newcert.pem` file into the CA Certificates list.
 - a. From the Management Console, select **Configuration > SSL > CA Certificates**.
 - b. Click **Import**; enter the certificate name in the CA Cert Name field.
 - c. Paste the certificate, being sure to include the `-----BEGIN CERTIFICATE-----` and the `-----END CERTIFICATE-----` statements in the `./bin/PEM/demoCA/private/CAcert` file.
 - d. Click **OK**.

Testing the Configuration

Import the root CA into your browser and construct an SSL interception policy.

Note: Detailed instructions on constructing an SSL interception policy are in ["Configuring SSL Rules through Policy"](#) on page 227.

You should not be prompted for any certificate warning.

Creating an Intermediate CA using Microsoft Server 2012 (Active Directory)

This section describes certificate management when creating an intermediate CA using Active Directory.

Before you begin:

- ❑ Verify the Windows 2012 system is an Active Directory server.
- ❑ Make sure IIS is installed on the server.
- ❑ Install the "Certificate Services" through the Server Manager. Enable **Active Directory Certificate Services** and select the **Certificate Authority mode** as **Enterprise root CA on the AD CS** (Active Directory Certificate Services).

All certificate management is done through the browser using the following URL:

`http://@ip_server/CertSrv`

For information on the following tasks, see:

- ❑ ["Install the root CA onto the browser:"](#) on page 252
- ❑ ["Create a ProxySG appliance keyring and certificate signing request:"](#) on page 252
- ❑ ["Sign the ProxySG appliance CSR:"](#) on page 252
- ❑ ["Import the subordinate CA certificate onto the ProxySG appliance:"](#) on page 253
- ❑ ["Test the configuration:"](#) on page 253

Install the root CA onto the browser:

1. Connect to `http://@ip_server/certsrv`.
2. Click **Download a CA Certificate, certificate chain, or CRL**.
3. Click **Install this CA Certificate**.

This installs the root CA onto the browser.

Create a ProxySG appliance keyring and certificate signing request:

1. From the Management Console, select the **Configuration > SSL > Keyrings** tab.
2. Create a new keyring. For detailed instructions on creating a new keyring, see ["Creating a Keyring"](#) on page 1119.
3. Create a Certificate Signing Request (CSR). For detailed instructions on creating a CSR, see ["Creating a CSR"](#) on page 1130.
4. To capture the CSR information, edit the keyring containing the CSR, and copy the **Certificate Signing Request** field content.
5. Click **Close**.

Sign the ProxySG appliance CSR:

1. Connect to `http://@ip_server/certsrv`.
2. Select **Request a certificate**.
3. Select **submit an advanced certificate request**.

4. On the next screen (**Submit a Certificate Request or Renewal Request**) paste the contents of the CSR into the **Base-64-encoded certificate request** field.
5. Select the Certificate Template **Subordinate Certification Authority**.
If this template does not exist, connect to the certificate manager tool on the Active Directory server and add the template.
6. Click **Submit**.
7. Download the certificate (not the chain) as **Base 64 encoded**.
8. Save this file on the workstation as `newcert.pem`.

Import the subordinate CA certificate onto the ProxySG appliance:

1. Open the file `newcert.pem` in a text editor and copy the contents, from the **BEGIN CERTIFICATE** through **END CERTIFICATE**; don't include any spaces after the dashes.
2. In the Management Console, select the **Configuration > SSL > SSL Keyrings** tab.
3. Select the keyring that has the CSR created; click **Edit**.

Note: Ensure this keyring is used as the issuer keyring for emulated certificates. Use policy or the SSL intercept setting in the Management Console or the CLI.

4. Click **Import** to paste the contents of the `newcert.pem` file. This imported the ProxySG appliance's subordinate CA certificate into the keyring.
5. To insure the ProxySG appliance trusts the newly added certificate, import the contents of the `newcert.pem` file into the CA Certificates list.
 - a. From the Management Console, select **Configuration > SSL > CA Certificates**.
 - b. Click **Import**; enter the certificate name in the **CA Cert Name** field.
 - c. Paste the certificate, being sure to include the `-----BEGIN CERTIFICATE-----` and the `-----END CERTIFICATE-----` statements in the `./bin/PEM/demoCA/private/CAcert` file.
 - d. Click **OK**.
 - e. Click **Apply**.

Test the configuration:

Import the root CA into your browser and construct an SSL interception policy. You should not be prompted for any certificate warning.

Note: Detailed instructions on constructing an SSL interception policy are in "[Configuring SSL Rules through Policy](#)" on page 227.

Chapter 10: Accelerating File Sharing

This chapter discusses file sharing optimization. File sharing uses the Common Internet File System (CIFS) protocol.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ ["About the CIFS Protocol"](#) on page 255
- ❑ ["About the Blue Coat CIFS Proxy Solution"](#) on page 256
- ❑ ["Configuring the ProxySG CIFS Proxy"](#) on page 258

About the CIFS Protocol

The CIFS protocol is based on the Server Message Block (SMB) protocol used for file sharing, printers, serial ports, and other communications. It is a client-server, request-response protocol. The CIFS protocol allows computers to share files and printers, supports authentication, and is popular in enterprises because it supports all Microsoft operating systems, clients, and servers.

File servers make file systems and other resources (printers, mailslots, named pipes, APIs) available to clients on the network. Clients have their own hard disks, but they can also access shared file systems and printers on the servers.

Clients connect to servers using TCP/IP. After establishing a connection, clients can send commands (SMBs) to the server that allows them to access shares, open files, read and write files—the same tasks as with any file system, but over the network.

CIFS is beneficial because it is generic and compatible with the way applications already share data on local disks and file servers. More than one client can access and update the same file, while not compromising file-sharing and locking schemes. However, the challenge for an enterprise is that CIFS communications are inefficient over low bandwidth lines or lines with high latency, such as in enterprise branch offices. This is because CIFS transmissions are broken into *blocks* of data; each block has a maximum size of 64 KB for SMBv1. When using SMBv1, the client must stop and wait for each block to arrive before requesting the next block. Each stop represents time lost instead of data sent. Therefore, users attempting to access, move, or modify documents experience substantial, work-prohibiting delays.

The second version of SMB (SMBv2) alleviates some of the inefficiencies in CIFS communication and improves performance over high latency links. Servers that support SMBv2 pipelining can send multiple requests/responses concurrently which improves performance of large file transfers over fast networks. While SMBv2 has some improvements, it does not address all of the performance issues of CIFS; for example, it cannot reduce payload data transferred over low bandwidth links.

About the Blue Coat CIFS Proxy Solution

The CIFS proxy on the ProxySG combines the benefits of the CIFS protocol with the abilities of the ProxySG to improve performance, reduce bandwidth, and apply basic policy checks. This solution is designed for branch office deployments because network administrators can consolidate their Windows file servers (at the data center) instead of spreading them across the network.




Figure 10–1 CIFS Proxy Traffic and Flow Diagram

Caching Behavior

The CIFS proxy caches the regions of files that are read or written by the client (partial caching) and applies to both read and write file activities. Also, the caching process respects file locking.

SMBv1 and SMBv2 share the same object cache, allowing a client using SMBv2 protocol to use objects cached by another client using SMBv1 (and vice versa). When SMBv2 protocol acceleration is disabled or the connection requires messages to be signed, the connection is placed into passthrough and object caching is not performed. However, the connection can still take advantage of byte caching and compression.

Note: Caching behavior can also be controlled with policy. See the *Content Policy Language Reference Guide* or the *Visual Policy Manager Reference Guide*.

Authentication

The CIFS proxy supports both server and proxy authentication in the following contexts.

Server Authentication

Permissions set by the origin content server (OCS) are always honored. Requests to open a file are forwarded to the OCS; if the OCS rejects the client access request, no content is served from the cache.

Note: NTLM/IWA authentication requires that the client knows what origin server it is connecting to so it can obtain the proper credentials from the domain controller.

Proxy Authentication

The ProxySG cannot issue a challenge to the user over CIFS, but it is able to make use of credentials acquired by other protocols if IP surrogates are enabled.

Policy Support

The CIFS proxy supports the `proxy`, `cache`, and `exception` policy layers. However, the SMB protocol can only return error numbers. Exception definitions in the forms of strings cannot be seen by an end user. Refer to the *Content Policy Language Reference* for supported CPL triggers and actions.

Access Logging

By default, the ProxySG uses a Blue Coat-derived CIFS access log format.

```
date time c-ip c-port r-ip r-port s-action s-ip cs-auth-group
cs-username x-client-connection-bytes x-server-connection-bytes
x-server-adn-connection-bytes x-cifs-method
x-cifs-client-read-operations x-cifs-client-write-operations
x-cifs-client-other-operations x-cifs-server-operations
x-cifs-error-code x-cifs-server x-cifs-share x-cifs-path
x-cifs-orig-path x-cifs-client-bytes-read x-cifs-server-bytes-read
x-cifs-bytes-written x-cifs-uid x-cifs-tid x-cifs-fid x-cifs-file-size
x-cifs-file-type x-cifs-fid-persistent
```

WCCP Support

If WCCP is deployed for transparency, you must configure WCCP to intercept TCP ports 139 and 445.

Configuring the ProxySG CIFS Proxy

This section contains the following sub-sections:

- ❑ ["About Windows Security Signatures" on page 258](#)
- ❑ ["Intercepting CIFS Services" on page 261](#)
- ❑ ["Configuring SMBv1 Options" on page 262](#)
- ❑ ["Configuring SMBv2 Options" on page 267](#)
- ❑ ["Reviewing CIFS Protocol Statistics" on page 268](#)

See Also

["About the CIFS Protocol" on page 255](#)

About Windows Security Signatures

Security signatures prevent the CIFS proxy from providing its full acceleration capabilities. Additionally, security signatures require a considerable amount of processing on both clients and servers. As their benefits are often superseded by link-layer security measures, such as VPNs and restricted network topology, the benefits are minimal and the drawbacks are high.

SMBv1

In order for the CIFS proxy to fully optimize SMBv1 traffic, the Windows clients cannot be configured with a requirement that security signatures always be used. The instructions for verifying this setting are detailed below.

In addition, if signing is required on the server, you must enable and configure SMB signing on the ADN concentrator. (See ["Enabling SMB Signing Support for SMBv1 Connections" on page 264.](#))

SMBv2

For SMBv2, if security signatures are always required on the client or the server, the CIFS proxy cannot fully optimize SMBv2 traffic. The proxy can perform byte caching and compression on this traffic, but it cannot perform object caching or protocol acceleration. If you want to fully optimize SMBv2 traffic, you must disable the setting that controls whether digital signing must always be used; this must be configured on clients and servers. If either side requires signing always be used, the SMBv2 connections will be passed through the proxy without full optimization.

If your clients are running Windows 8, you can optionally disable Secure Dialect Negotiation on clients for better performance; see "[Disable Secure Dialect Negotiation](#)" on page 261.


Verify Security Signature Settings in Windows

By default, the **Digitally sign communications (always)** setting is disabled in Windows, except in the case of a Domain Controller installation. If this setting has been enabled, you will need to disable it in order to fully optimize CIFS traffic.

To verify the security signature settings in Windows:

Note: This procedure follows the Control Panel Classic View format. The screen shots represent Microsoft Windows XP.


1. In each Windows client, select **Start > Control Panel > Administrative Tools > Local Security Policy**. The Local Security Settings dialog appears.



2. Select **Local Policies > Security Options**.

3. Right-click **Microsoft network client: Digitally sign communications (always)** and select **Properties**. A configuration dialog appears.

Note: In Windows 2000, this option is called **Digitally sign client communications (always)**.



4. Select **Disabled**. Click **Apply** and **OK**.
5. Close all Control Panel dialogs.

Important: If the server is an ADS/Domain controller, you must set the same security settings for both **Administrative Tools > Domain Controller Security Policy** and **Administrative Tools - > Domain Security Policy**. Otherwise, you cannot open file shares and Group Policy snap-ins on your server.

6. You must reboot the client to apply this configuration change.
7. *SMBv2 only:* Repeat these steps on the servers you want the CIFS proxy to optimize. On the server, the option is called **Microsoft network server: Digitally sign communications (always)**.
8. *SMBv1 only:* If the server requires signing, enable and configure SMB signing on the ADN concentrator. See "[Enabling SMB Signing Support for SMBv1 Connections](#)" on page 264.

Disable Secure Dialect Negotiation

Secure Dialect Negotiation allows servers and clients to detect an attacker's attempts to eavesdrop on server-client communication and downgrade the negotiated SMBv2 protocol dialect.

If your clients run Windows 8, you can optionally disable Secure Dialect Negotiation to improve CIFS performance; however, CIFS operations still function correctly if the feature is enabled.

Disable Secure Dialect Negotiation on the client:

1. In Windows, open the Registry Editor.
2. Look for the following key:
`HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\LanmanServer\Parameters`
3. Select the key, and then add a new DWORD value with the name `RequireSecureNegotiate` and value data 0.

Intercepting CIFS Services

By default (upon upgrade and on new systems), the ProxySG has CIFS services configured for transparent connections on ports 139 and 445. Blue Coat creates listener services on both ports because different Windows operating systems (older versus newer) attempt to connect using 139 or 445. For example, Windows NT and earlier only used 139, but Windows 2000 and later try both 139 and 445. Therefore only configuring one port can potentially cause only a portion of Windows 2000 and newer CIFS traffic to go through the proxy.

A transparent connection is the only supported method; the CIFS protocol does not support explicit connections.


Also, by default these services are configured to accept all IP addresses in **Bypass** mode. The procedure in this section describes how to change them to **Intercept** mode, and explains other attributes within the service.

Adding and Configuring New CIFS Services

If you require a CIFS service to intercept a port other than the default 139/445 ports, you can create a new service (and specify a default or custom service group). This general procedure is described in "[Creating Custom Proxy Services](#)" on page 120.

To configure the CIFS proxy to intercept file sharing traffic:

1. From the Management Console, select **Configuration > Services > Proxy Services**.



2. Intercept CIFS traffic:

- a. Scroll the list of service groups, click **Standard**, and click **CIFS** to expand the CIFS services list.
- b. Notice the **Action** for each default service (ports 139 and 445) is **Bypass**. Select **Intercept** from the drop-down list(s).

3. Click **Apply**.


Now that the ProxySG is intercepting CIFS traffic, configure the CIFS proxy options for SMBv1 ("[Configuring SMBv1 Options](#)" on page 262) or SMBv2 ("[Configuring SMBv2 Options](#)" on page 267).

Configuring SMBv1 Options

When using SMBv1, you can configure options for file reading/writing, remote storage optimization, and folder management and caching. This section describes these options and why they might require changing based on your branch deployment.

To view/change the SMBv1 configuration options:

1. In the Management Console, select the **Configuration > Proxy Settings > CIFS Proxy > SMBv1** tab.
2. To accelerate SMBv1 connections, make sure the **Enable protocol acceleration for SMBv1 connections** check box is selected.



3. Configure the SMBv1 options:

- a. **Read Ahead:** The appliance attempts to *anticipate* what data might be requested next, fetches it, and caches it; this reduces the latency of the connection. The ProxySG might partially cache a requested object (the part directly requested and viewed by the client). Enabled by default.

If applications frequently perform large amounts of non-sequential file access, disable **Read Ahead** to reduce the amount of unnecessary data being fetched into the cache.

- b. **Write Back:** This setting applies when clients attempt to write to a file on the core server. Without write back, a client would experience substantial latency as it sends data chunks and waits for the acknowledgement from the server to send subsequent data chunks. With **Write Back** set to **Full**, the branch ProxySG sends acknowledgements to the client, prompting the client to send subsequent data without waiting for an acknowledgement from the core server. Meanwhile, the ProxySG forwards the data from the client to the core server through the compressed TCP connection.

For best performance, set **Write Back** to **Full** (default).

- c. **Remote Storage Optimization:** When this option is enabled, Windows Explorer modifies the icons of uncached folders on remote servers, indicating to users that the contents of the folder have not yet been cached by the ProxySG. Disabled by default.

When remote storage optimization is enabled, the ProxySG reports to the client that files are *offline* if the file is not in cache. This is designed to reduce the amount of chatter that a client will generate for files. By default, Windows Explorer does not show offline files in the search results. In order to force Windows Explorer to show offline files, you can select **Search tape backup** in the Windows Explorer advanced search options.

- d. **Suppress Folder Customization:** When this option is enabled, remote folders are displayed in the default view, without any view customizations (such as showing thumbnails instead of icons). This setting speeds the display of remote folders, especially on slow links. Disabled by default.

- e. **Never Serve Directories After Expiration:** When this option is enabled and **Directory Cache Time** is past its expiration, directories are refreshed synchronously instead of in the background. This is needed when the set of visible objects in a directory returned by a server can vary between users. Disabled by default.
 - f. **Directory Cache Time:** This option determines how long directory information remains in the object cache. Changes made to a directory by clients not using the ProxySG are not visible to ProxySG clients if they occur within this time interval. The default cache time is 1 minute. Blue Coat recommends keeping this value low to ensure clients have access to the most current directory information; however, you can set it longer if your applications use CIFS to access files. For example, the cache responds faster if it knows directory *x* does not contain the file and so moves on to directory *y*, which reduces the number of round trips to the file server.
- 4. Click **Apply** to save your settings.
 - 5. To configure SMB signing, see ["Enabling SMB Signing Support for SMBv1 Connections"](#) below.
 - 6. To configure SMBv2, see ["Configuring SMBv2 Options"](#) on page 267.

Enabling SMB Signing Support for SMBv1 Connections

SMB signing is a Microsoft-devised security mechanism that attempts to prevent *man-in-the-middle* attacks. If a network administrator configures SMB signing on clients and servers, signatures are added to the packet header. A decrypted signature by the recipient server or client indicates a valid packet. If the signature is malformed or not present, or if the SMB packet is compromised, the client or server rejects and drops the packet.

The administrator can configure SMB signing in one of two modes:

- ❑ **Enabled**—Clients that support SMB signing connect to SMB-signing enabled servers with signed SMB sessions. Clients that do not support SMB signing are also able to connect to SMB-signed servers, but the SMB sessions are not signed). By default, SMB signing is enabled for outgoing SMB sessions on Windows 7, Windows 2008, Windows Vista, Windows XP, Windows 2000, Windows NT4.0 and Windows 98 operating systems.
- ❑ **Required Always**—The client or server is only able to send or receive to counterparts that support SMB signing. Because this limits the systems to which a client or server can communicate, SMB signing is not commonly configured as always required. However, it is required for incoming SMB sessions on Windows Server 2003/2008-based domain controllers which means that if the domain controller also acts as a file server, sessions with the SMB signing-enabled clients listed in the previous bullet are signed.

Because the ProxySG potentially resides between SMB signing-configured clients and servers, it must be able to provide file sharing (CIFS proxy) acceleration and optimization without compromising SMB signing. To achieve this, the ProxySG

serves as a *virtual user* (SMB signing is transparent to users) when the option to optimize SMB-signed traffic is enabled. What occurs depends on the configuration of the OCS.




Figure 10–2 OCS Configuration Determines ProxySG Process

Traffic between the branch and the Concentrator is not signed. Regardless of the OCS SMB configuration, the client receives a message that the packets do not require SMB signatures (see Figure 10–2 above). This enables the ProxySG to intercept the CIFS protocol and provide optimization. Because of slightly higher use of the CPU, enabling SMB signing on clients and servers slightly decreases performance.

Notes

- ❑ SMB signing is not supported for SMBv2 connections on the ProxySG.
- ❑ If an error occurs, such as problems with the specified domain access credentials, the ProxySG allows the traffic to pass through.


Prerequisites

- ❑ Before configuring SMB signing on the ProxySG, you must create a user in the domain that represents the ProxySG. When SMB signing is always required by the OCS, the ProxySG CIFS proxy uses this virtual user's credentials. This user *cannot* be a guest or anonymous.

- ❑ The Windows clients cannot be configured to always require signing. See ["About Windows Security Signatures"](#) on page 258.

To enable SMB-signed packet optimization for SMBv1 connections when the server requires signing:

1. From the Management Console of the Concentrator ProxySG (not the branch), select **Configuration > Proxy Settings > CIFS Proxy > SMBv1**.



2. In the **SMB Signing** area, select **Enable protocol optimizations on signed SMB traffic using the following credentials**.
3. In the **Username** field, enter the user name that you created in the domain. Ensure you enter the name exactly as created. It is optional to enter the **Domain** to which the username belongs.
4. Enter the username password that the ProxySG sends to access the domain:
 - a. Click **Set password**. The Set Password dialog displays.
 - b. Enter the password in both fields.
 - c. Click **OK**.
5. Click **Apply**.

SMB Signing Log Entries and Statistics

When SMB signing optimization is enabled, the following messages are possible:

- ❑ **Event Log**—Entry when authentication fails because of a problem with SMB access credentials. This entry marks the first occurrence; subsequent occurrences are not entered.
- ❑ **Debug Logs**—
 - Signature computation fails.
 - Packet signing error.

- Each time a domain or server authentication result, success or failure, occurs with pass-through connections.
- ❑ Active Sessions—Authentication failure when using SMB user credentials during a pass-through connection.
- ❑ Access Log—Successful logging on and off using SMB user credentials.

Configuring SMBv2 Options

Note: SMBv2 support in an ADN deployment requires both the branch and concentrator peers to be running SGOS 6.4 or higher. If they aren't, SMBv2 connections are downgraded to SMBv1.


The CIFS proxy supports SMBv2 protocol enhancements including data pipelining, request compounding, larger reads and writes, improved scalability for file sharing, durable opens during temporary loss of network connectivity, and the leasing mechanism for caching. Note that protocol optimization cannot be applied to SMBv2 connections that require messages to always be signed.

Compatibility

- ❑ The following Microsoft client and server operating systems: Windows Vista, Windows 7, Windows Server 2008, Windows Server 2008 R2.
- ❑ Samba SMB clients and servers
- ❑ EMC filers
- ❑ NetApp filers

To view/change the SMBv2 configuration options:

1. In the Management Console, select the **Configuration > Proxy Settings > CIFS Proxy > SMBv2** tab.



2. Choose one of the following ways to handle SMBv2 connections:
 - **Enable protocol acceleration for SMBv2 connections**—Unsigned SMBv2 connections are accelerated with object caching and any other ADN optimizations that are enabled. SMBv2 connections that require signing are passed through, allowing the proxy to accelerate them with byte caching and compression techniques (if enabled). No object caching is performed on the pass-through connections.

- **Downgrade SMBv2 connections to SMBv1**—Attempts to force the negotiation of SMBv1 for the connection. If downgrading isn't possible (for example, if the client negotiates SMBv2 directly or if SMBv1 protocol acceleration is disabled), the connection is passed through, allowing it to be accelerated with any ADN optimizations that are enabled for the CIFS service. No object caching is performed on the pass-through connections.
- **Disable protocol acceleration for SMBv2 connections**—All SMBv2 connections are passed through, allowing the proxy to accelerate them with any ADN optimizations that are enabled for the CIFS service. No object caching is performed on SMBv2 connections.

3. Click **Apply**.

Enabling CIFS Access Logging


By default, the ProxySG is configured to use the Blue Coat CIFS access log format. Enable Access Logging on the **Configuration > Access Logging > General** page.

Reviewing CIFS Protocol Statistics

After CIFS traffic begins to flow through the ProxySG, you can review the statistics page and monitor results in various CIFS categories. The presented statistics are representative of the client perspective.

To review CIFS statistics:

1. From the Management Console, select **Statistics > Protocol Details > CIFS History**.



2. View statistics:
 - a. From the **Service** or **Proxy** drop-down list, select **CIFS**.
 - b. Select a statistic category tab:
 - **CIFS Objects:** The total number of CIFS-related objects processed by the ProxySG (read and written).
 - **CIFS Bytes Read:** The total number of bytes read by CIFS clients.
 - **CIFS Bytes Written:** The total number of bytes written by CIFS clients (such as updating existing files on servers).
 - **CIFS Clients:** The total number of connected CIFS clients.
 - **CIFS Bandwidth Gain:** The total bandwidth usage for clients (yellow) and servers (blue), plus the percentage gain.
 - c. The graphs display three time metrics: the previous 60 minutes, the previous 24 hours, and the previous 30 days. Select **Duration:** from the drop-down list. Roll the mouse over any colored bar to view details.
3. (Optional) You can change the scale of the graph to display the percentage of bar peaks to display.

Chapter 11: Managing Outlook365 Applications

This chapter discusses intercepting Skype for Business, as well as the Endpoint Mapper service and MAPI proxy, which function together to intercept traffic generated by Microsoft Outlook clients and accelerate traffic over the WAN.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ [Section A: "The Outlook Proxies"](#) on page 271
- ❑ [Section B: "Endpoint Mapper and MAPI Configuration"](#) on page 278

Section A: The Outlook Proxies

This section discusses the Endpoint Mapper and MAPI proxies and how they work together to accelerate Outlook email traffic.

- ❑ ["About the Endpoint Mapper Proxy"](#) on page 271
- ❑ ["About the MAPI Proxy"](#) on page 272
- ❑ ["About MAPI Over HTTP"](#) on page 275

About the Endpoint Mapper Proxy

The Endpoint Mapper proxy is a key component of Symantec's solution for accelerating Outlook email traffic. Endpoint Mapper is a Remote Procedure Call (RPC) service that allows communication between Outlook clients and Exchange servers. As an RPC client, Outlook sends a message to Endpoint Mapper, asking what port Exchange is listening on; then Outlook uses the supplied port to communicate with the server.

The challenges occur when these communications occur between Outlook clients at branch offices and Exchange servers located in core locations. The user experience is poor because of low available bandwidth or high latency lines. This is where the Endpoint Mapper proxy can help.

This proxy intercepts the RPC client request for a particular RPC service. When the RPC client connects to the service, the Endpoint Mapper proxy secondary service intercepts the request and tunnels it. Substantial performance increase occurs because:

- ❑ The ProxySG caches server information, negating the requirement to connect to an upstream server for repeated requests.
- ❑ The ProxySG at the branch office (the *Branch peer*) compresses RPC traffic and sends it over the TCP connection to the ProxySG at the core (the *Concentrator peer*), which decompresses the data before sending it to the RPC server.

The Endpoint Mapper proxy can be deployed in both transparent and explicit modes. Intercepting RPC traffic is part of the complete solution that includes the MAPI proxy.

Note: Only Microsoft RPC version 5.0 is supported. If the RPC version is not 5.0, the connection is terminated.

About the MAPI Proxy

MAPI is the protocol used by Microsoft Outlook (client) to communicate with Microsoft Exchange (server), most commonly for e-mail applications. MAPI itself is based on the Microsoft Remote Procedure Call (RPC).

Because MAPI is based on RPC, it suffers from the same performance inherent with RPC communications. As enterprises continue to trend toward consolidating servers, which requires more WAN deployments (branch and remote locations), e-mail application users experience debilitating response times for not only sending and receiving mail, but accessing message folders or changing calendar elements.

With the release of Exchange Server 2003 and subsequent versions of Outlook, Microsoft introduced data encoding to enhance the efficiency and security of file transfers. However, file encoding prevents data sent with the MAPI protocol from matching with data sent using *other* protocols (HTTP, CIFS, FTP, etc.), thereby limiting byte cache effectiveness.

About the Blue Coat MAPI Solution

The MAPI proxy is similar to and actually works in conjunction with the Endpoint Mapper proxy in that it intercepts and accelerates RPCs; however, MAPI is always deployed transparently and does not listen on a specific port or port range. Instead, when configured to do so, the Endpoint Mapper proxy *hands off* Outlook/Exchange traffic to the MAPI proxy (but the Endpoint Mapper proxy functionality is still required to make an RPC connection).

The MAPI proxy itself is a *split proxy*, which is only viable in a deployment that consists of a ProxySG at the branch office and a Concentrator ProxySG at the core. A split proxy employs co-operative processing at the branch and the core to implement functionality that is not possible in a standalone proxy. In the case of the MAPI proxy, cooperation exists between the ProxySG appliances at the branch and the core to reduce the number of RPCs sent across the WAN. The TCP connection between the Branch and Concentrator peers makes use of byte caching for acceleration.

MAPI compression includes all files and supported protocols sent from Microsoft Outlook. It also improves general performance, bandwidth and, in certain cases, application-level latency.

In summary, the Symantec MAPI solution supports the following acceleration techniques:

- Protocol optimizations

- ❑ Byte caching
- ❑ Compression
- ❑ Upload/download optimizations

The following diagram illustrates a typical MAPI communication flow:




Figure 11–1 MAPI Proxy Deployment and Flow Diagram

Reducing RPC Messages Across the WAN

The MAPI proxy batching feature reduces the number of RPC messages traversing the WAN during attachment download and upload.

- ❑ **Attachment download optimization** If the protocol and Exchange version permit, the Concentrator peer will either batch attachments that have multiple simultaneous RPC requests or request larger data chunks than the Outlook client requested. The Concentrator peer does attachment data read ahead and forwards it to the Branch, so that once Outlook requests the next data chunk, the Branch peer already has it available.
- ❑ **Attachment upload optimization** The Branch peer simulates the Exchange server by generating the attachment data acceptance response locally; this allows Outlook to send the next data fragment, thereby reducing the response round-trip time over the WAN, which saves time and bandwidth.

Maximizing Cross Protocol Byte-Cache Hits

The Symantec MAPI compression handling feature allows data encoded (or compressed) by Microsoft Outlook and Exchange to be byte cached and thereby accelerated. This feature improves bandwidth, especially when sending and receiving large attachments using Microsoft Outlook.

For example, when a user sends an e-mail with an attachment, Outlook encodes the data to the Exchange server. As the e-mail is sent across the line, the Branch peer intercepts and decodes the attachment data. Because the Branch peer sends the data across the WAN in a plain format, it can be byte-cached with all other supported protocols (CIFS, HTTP, FTP, etc.), thereby increasing cross-protocol hits. After the data reaches the concentrator ProxySG, it is encoded back to the Outlook standard and processed by the Exchange server.

When a user makes a receive request, the concentrator ProxySG decodes the data from the Exchange server. After the data reaches the Branch peer, it is once again encoded to the original format and processed by the Outlook client.

Currently, MAPI compression handling supports improved byte caching for MAPI 2000/2003. Both the Branch and Concentrator peers must run the same version of SGOS for MAPI compression functionality.

Note: Attachments sent using MAPI compression are transferred in plain over WAN when secure ADN is not used. Branch to Outlook and Concentrator to Exchange data is obfuscated using the native Microsoft encoding format.

Maintaining Exchange Connections

The MAPI proxy Keep-Alive feature allows the ProxySG to maintain the connection to the Exchange server after the user has logged off from Outlook. Determined by the configurable interval, the MAPI proxy checks the Exchange server for new mail. ADN Optimization allows the connection to remain warm so that when the user logs on again to Outlook, the number of retrieved bytes is lower, which provides better performance.

The MAPI proxy remembers each user that is logged on or off. If the duration exceeds the specified limit, or when the user logs back into the mail application, the Keep-Alive connection is dropped.

Supported Microsoft Outlook Clients and Exchange Servers

Refer to the following table to determine which MAPI protocol is supported if you are using a specific Exchange and Outlook combination.

Table 11–1 Supported ProxySG Exchange/Outlook Servers

	Exchange 2003	Exchange 2007	Exchange 2010*	Exchange 2013*	Exchange 2016*
Outlook 2003	MAPI 2003	MAPI 2003	MAPI 2003	MAPI 2003	MAPI 2003
Outlook 2007*	MAPI 2003	MAPI 2007	MAPI 2007	MAPI 2007	MAPI 2007
Outlook 2010*	MAPI 2003	MAPI 2007	MAPI 2010	MAPI 2010	MAPI 2010
Outlook 2013*	MAPI 2003	MAPI 2007	MAPI 2010	MAPI 2013	MAPI 2013
Outlook 2016*	MAPI 2003	MAPI 2007	MAPI 2010	MAPI 2013	MAPI 2016
*MAPI encryption enabled by default					

MAPI Backward Compatibility

SGOS allows MAPI backward compatibility, allowing functionality during upgrade/downgrade cycles and other instances when the appliances at the branch office and core are running different versions. As a result, any ongoing changes to the ProxySG appliances will not break application usability.

When the Branch and Concentrator peers encounter a MAPI version mismatch, they negotiate down to the lowest common version. Depending on which version of MAPI has been negotiated to, certain features found in later versions will not function.

For example, if the Branch peer runs SGOS 5.3 and the Concentrator peer runs SGOS 5.4, they will negotiate to SGOS 5.3. Because SGOS 5.3 does not support MAPI compression, users will not benefit from cross protocol byte-cache hits with CIFS or other compatible protocols.

Note: A warning appears in the Active Sessions at the branch office when connections are affected by a version downgrade.

About MAPI Over HTTP


MAPI over HTTP tunnels traffic over an HTTPS connection and accepts connections from the HTTP proxy instead of from the Endpoint Mapper Proxy. This protocol was introduced in Microsoft Outlook 2013 SP1 and it replaces RPC over HTTP.

When using this protocol with a ProxySG appliance, the appliance removes MAPI's compression from traffic sent over an ADN and applies its own compression instead.

About Encrypted MAPI

This feature provides the ability to transparently accelerate encrypted MAPI traffic between the Outlook client and the Exchange server. The ability to decrypt and encrypt MAPI is transparent to the user, with no knowledge of the user's password.

This feature assumes your ADN network is set up as follows.



The encrypted MAPI acceleration feature expects the Outlook client to use the Simple and Protected Negotiation (SPNEGO) security protocol, and as a result the proxy will negotiate NTLM protocol on the client side and Kerberos on the server side. SPNEGO is used when a client application wants to authenticate to a remote server, but neither end is sure what authentication protocols the other supports.

For configuration details, see ["Optimizing Encrypted MAPI Traffic"](#) on page 282.

Encrypted MAPI Requirements

The ProxySG encrypted MAPI feature has the following requirements:

- ❑ ADN must be configured with at least one Branch and one Concentrator peer. The peers must be running SGOS 6.2 or later and be configured to use an SSL device profile and secure ADN.
- ❑ An SSL license is required for secure ADN on the Branch and the Concentrator peers.
- ❑ The Outlook clients must be configured to use Kerberos/NTLM Password Authentication (Outlook 2003) or Negotiate Authentication (Outlook 2007, Outlook 2010) logon network security. The Exchange server must be enabled to support Kerberos security protocol and the Domain Controller must be enabled to support both Kerberos and NTLM LAN authentication protocols.
- ❑ The clocks on the Branch and Concentrator peers must be synchronized with the Domain Controller clock.

- ❑ The Branch peer must be joined to each Windows domain to which your Exchange server(s) and Outlook users belong. For example, if users are created in domain A and the Exchange server resides in domain B (which has a trust relationship with domain A), the ProxySG must be joined to both domains.
- ❑ The Branch peer must be configured to be trusted for delegation for exchangeMDB services and must act as an Active Directory member host.

Encrypted MAPI Limitations

The encrypted MAPI feature has the following limitations on the ProxySG:

- ❑ The encrypted MAPI solution on the ProxySG does not support batching.
- ❑ Encrypted MAPI 2000 is not supported on the ProxySG.
- ❑ Non-secure ADN can be reported in the Active Sessions at the branch even though secure ADN is enabled on the Branch and Concentrator peers. This can happen when Outlook establishes a plain connection with the Exchange server and then switches to the secure authentication level in the middle of a MAPI conversation. When this happens, the encrypted MAPI session goes through a plain ADN tunnel, without acceleration benefits.

To prevent this, enable the **Secure all ADN routing and tunnel connections** option.

- ❑ Encrypted MAPI is not supported if the Branch peer fails to authenticate the user by using NTLM and Kerberos authentication protocols within the Exchange domain.

Section B: Endpoint Mapper and MAPI Configuration

This section discusses the following configuration topics:

- ❑ "Configuring the Endpoint Mapper Service"
- ❑ "Using the MAPI Proxy" on page 279
- ❑ "Optimizing Encrypted MAPI Traffic" on page 282


Configuring the Endpoint Mapper Service

By default (upon upgrade and on new systems), the ProxySG has an Endpoint Mapper service configured on port 135. The service is configured to listen to all IP addresses, but might be set in **Bypass** mode (depending on the initial configuration performed by a network administrator).

In order to manage Outlook traffic, the Endpoint Mapper service must be intercepted.

To set the Endpoint Mapper service to intercept:

1. From the Management Console, select **Configuration > Services > Proxy Services**.



2. Change the Endpoint Mapper service to intercept:
 - a. Scroll the list of service groups, click **Standard**, and select **Endpoint Mapper**.
 - b. If the **Action** for the default service (port 135) is set to **Bypass**, select **Intercept** from the drop-down list(s).
3. Click **Apply**.

Adding a New Endpoint Mapper Service

The ProxySG allows you to add new Endpoint Mapper services. Consider the following scenario: you want the ProxySG to exclude (bypass) an IP address/subnet from MAPI acceleration because that network segment is undergoing routine maintenance. To learn more about adding custom services, see ["Creating Custom Proxy Services"](#) on page 120.

Bypassing Endpoint Mapper Traffic

Certain scenarios might require you to change the Endpoint Mapper service from **Intercept** to **Bypass**. For example, you need to take an Endpoint Mapper service offline for maintenance. When an Endpoint Mapper changes from Intercept to Bypass, the ProxySG closes not only the primary connections (such as connections to a Microsoft Exchange server on port 135), but also the secondary connections, which are used to intercept further RPC requests on mapped ports. The result is fully bypassed Endpoint Mapper traffic.

Reviewing Endpoint Mapper Proxy Statistics

After RPC traffic begins to flow through the ProxySG, you can review the statistics page and monitor results in various categories. The presented statistics are representative of the client perspective.

Management Console Statistics Pages

Endpoint Mapper statistics display across multiple pages:

- ❑ **Statistics > Traffic Mix** tab—Service and proxy data; bandwidth use and gain; client, server, and bypassed bytes. Includes all traffic types, but you can limit the scope to Endpoint Mapper data.
- ❑ **Statistics > Traffic History** tab—Service and proxy data; bandwidth use and gain; client, server, and bypassed bytes. Select Endpoint Mapper service or proxy (related to MAPI, as described in ["Configuring the MAPI Proxy"](#) on page 280).
- ❑ **Statistics > Active Sessions**—The **Proxied Sessions** and **Bypassed Connections** tabs display statistics filtered by various criteria, such as port or service type (select **Endpoint Mapper**).

Statistic URL Pages

Endpoint Mapper proxy statistics pages are viewable from Management Console URLs. This page displays various, more granular connection and byte statistics.

`https://SG_IP_address:8082/epmapper/statistics`

Using the MAPI Proxy

This section discusses the following topics:

- ❑ ["Configuring the MAPI Proxy"](#) on page 280
- ❑ ["Reviewing MAPI Statistics"](#) on page 281

Configuring the MAPI Proxy


This section discusses how to configure the MAPI proxy acceleration features.

For more information, see the following sections:

- "About the MAPI Proxy" on page 272
- "Reviewing MAPI Statistics" on page 281

To view/change the MAPI Proxy configuration options:

1. In the Management Console, select **Configuration > Proxy Settings > MAPI Proxy**.



2. Configure the MAPI proxy configuration options:
 - a. **Enable Endpoint Mapper to MAPI Handoff:** Use this option to enable MAPI acceleration. The Endpoint Mapper proxy sends Microsoft Outlook and Exchange RPC communications to the MAPI proxy, which is used to manage the data. The routing connections from the branch to the core remains under the control of the Endpoint Mapper service.

Note: A secondary TCP connection is created to handle all non-MAPI traffic. No changes to the Endpoint Mapper service or proxy are required.

- b. **Enable acceleration for encrypted MAPI:** Select this option if you want to accelerate encrypted MAPI traffic. To use this option you must join the appliance to each Windows domain to which your Exchange server belongs and Outlook users are created. You must then select the **Domain alias** that is associated with that domain to enable encrypted MAPI acceleration. If you do not select a **Domain alias**, the appliance

will bypass encrypted MAPI traffic (and the associated traffic will show the `Domain alias not set` message in Active Sessions). If you have not yet joined the appliance to a Windows domain, see ["Integrate the ProxySG Appliance into the Windows Domain"](#) on page 1007 for instructions.

Note: Before enabling acceleration for encrypted MAPI, make sure you have performed the required setup tasks on the Domain Controller, and on the Branch and Concentrator peers. See ["Optimizing Encrypted MAPI Traffic"](#) on page 282 for details.

- c. **Enable batching of attachment uploads/downloads:** If enabled, this option reduces the MAPI message count sent over the ADN tunnel during attachment upload and download. This reduction in message roundtrips saves time.

Note: For the batching option to produce additional time gains, the **Cached Exchange Mode** option on the Outlook client must be disabled.

- d. **Enable keep-alive for disconnected clients:** After a user closes Outlook, the MAPI RPC connection remains and the ProxySG continues to receive incoming messages to this account. If disabled (the default), no attempts to contact the server occur until the next time the user logs into his/her Outlook account. This might create a noticeable decrease in performance, as the queue of unreceived mail is processed.
- **Interval:** How often the MAPI proxy contacts the Exchange server to check for new messages.
 - **Duration:** How long the MAPI proxy maintains the connection to the Exchange server. The connection is dropped if the duration exceeds this value or once a user logs back in to the mail application.
 - **Maximum Sessions:** Limits the number of occurring active keep-alive sessions. If a new keep-alive session starts, and the specified limit is already exceeded, the oldest keep-alive session is *not* dropped but no new keep-alive sessions are created.
3. Click **OK**.
4. Click **Apply**.

Reviewing MAPI Statistics

After MAPI traffic begins to flow through the ProxySG, you can review the statistics page and monitor results in various MAPI categories. The presented statistics are representative of the client perspective.

To review MAPI History:

1. From the Management Console, select **Statistics > MAPI History**.

2. View statistics:
 - a. Select a statistic category tab:
 - **MAPI Clients Bytes Read:** The total number of bytes read by MAPI clients.
 - **MAPI Clients Bytes Written:** The total number of bytes written by MAPI clients.
 - **MAPI Clients:** The total number of MAPI connections.
 - b. The graphs display three time metrics: the previous 60 minutes, the previous 24 hours, and the previous month. Roll the mouse over any colored bar to view the exact metric.
3. (Optional) You can change the scale of the graph to display the percentage of bar peaks to display.

To review MAPI Active Sessions:

1. From the Management Console, select the **Statistics > Active Sessions > Proxied Sessions** tab.
2. From the first **Filter** drop-down list, select **Proxy**; from the second drop-down list, select **MAPI**.
3. Click **Show**. The **Proxied Sessions** area displays MAPI statistics.

Optimizing Encrypted MAPI Traffic

Enabling optimization of the encrypted MAPI protocol requires the following tasks. If these tasks are not performed, the ProxySG tunnels MAPI traffic without optimization. Some of these tasks are performed on the Domain Controller, some on the Branch peer, and others on the Concentrator peer.

Task #	Task	Reference
1	Prepare the Domain Controller to support the Trust Delegation feature.	"Prepare the Domain Controller to Support Trust Delegation" on page 283
2	Ensure that the clocks on the ProxySG appliances at the branch office and core are synchronized with the Domain Controller.	"Synchronize the ProxySG Appliances and DC Clocks" on page 283
3	Configure secure ADN between the Branch and Concentrator peers.	"Verify Secure ADN" on page 284
4	Join the ProxySG at the branch to the primary domain (the same domain where the Exchange server is installed).	"Join the Branch Peer to the Primary Domain" on page 285
5	On the Domain Controller, configure Trust Delegation for the host name of the ProxySG at the branch office.	"Configure the Domain Controller to Trust the ProxySG Host" on page 285

Task #	Task	Reference
6	Enable MAPI encryption on the ProxySG at the branch office.	"Enable MAPI Encryption Support" on page 286


Prepare the Domain Controller to Support Trust Delegation

Note: Only the Primary Domain Controller requires the new configuration; the configuration automatically replicates to the Backup Domain Controller.

The trust delegation feature (configured in a later task) requires that the domain functional level be at Windows Server 2003 (or newer).

If you need to raise the functional level:

1. On the Domain Controller, select **Administrative Tools**, and open **Active Directory Domains and Trusts**.
2. Right-click the domain and select **Raise Domain Functional Level**.



3. From **Select an available domain functional level**, select **Windows Server 2003** (or newer) and click **Raise**.

Note: After raising the domain functional level to Windows Server 2003 from Windows 2000, you cannot add additional Windows 2000 servers to this domain.

Synchronize the ProxySG Appliances and DC Clocks

The clocks on the ProxySG appliances at the branch office and core must be synchronized with the clock on the Domain Controller. Note that a Branch peer cannot join an AD domain unless its internal clock is in sync with the Domain Controller. In addition, if the Concentrator is out of sync with the other clocks, it will not be able to establish an encrypted MAPI session.

To ensure that the ProxySG clocks are synchronized with the Domain Controller clock, use either of the following techniques:


- ❑ Specify the same NTP servers for the ProxySG appliances and the Domain Controller.
- ❑ Configure the ProxySG appliances to use the Domain Controller as the NTP source server.

ProxySG NTP configuration options are located on the **Configuration > General > Clock** tab.


Verify Secure ADN

The Branch and Concentrator peers must have SSL licenses and be configured to use the same SSL device profile and secure ADN.

Configuring the ProxySG appliances for Secure ADN



1. On the Branch peer, select the **Configuration > ADN > General > Device Security** tab.
2. Verify an **SSL Device Profile** is selected; if not, select one (if you need to create one, refer to the Help System).
3. Click **Apply** to commit any changes.



4. Select the **Configuration > ADN > General > Connection Security** tab.
5. In the **Secure-Outbound Mode** area, verify a secure option is selected.
6. Click **Apply** to commit any changes.

Join the Branch Peer to the Primary Domain

One of the requirements for accelerating encrypted MAPI traffic is that the ProxySG at the branch office must be joined to each Windows domain to which your Exchange server(s) and Outlook users belong. For example, if users are created in domain A and the Exchange server resides in domain B (which has a trust relationship with domain A), the ProxySG must be joined to both domains.


For details on how to join the domain, see ["Join the ProxySG Appliance to the Windows Domain"](#) on page 1008.

Configure the Domain Controller to Trust the ProxySG Host

For the ProxySG to be able to authenticate Exchange users, the Domain Controller must trust the ProxySG host for delegation. Note that the ProxySG host can be trusted to delegate for multiple Exchange servers.

Trusting the ProxySG as a Host

1. On the Domain Controller, select **Administrative Tools**, and open **Active Directory Users and Computers**.



2. Under **DomainName/Computers**, double-click the ProxySG host to display the Properties dialog.

- a. On the **Delegation** tab, click **Trust this computer for delegation to specified services only**.

If you don't see the **Delegation** tab, you did not raise the delegation level to Windows Server 2003 or newer. See ["Prepare the Domain Controller to Support Trust Delegation"](#) on page 283.


- b. Click **Use any authentication protocol**.
- c. Click **Add**; in **Add Services**, click **Users and Computers**.
- d. In the **Enter the object names to select (examples)** field, enter the name of the Exchange server for which the system will be trusted to delegate and click **OK**.
- e. In **Add Services**, click the **Exchange MDB** that will be trusted for delegation and click **OK**.
- f. Repeat steps d and e for any other endpoint Exchange servers that accept MAPI connections.

- g. Click **OK** to close the Properties dialog.

Enable MAPI Encryption Support

After completing the previous preparatory tasks, you are now ready to configure the Branch peer to intercept and optimize encrypted MAPI traffic. This setting is enabled by default on fresh installations; it is disabled on upgraded systems.

Enabling MAPI Encryption Support



1. In the Management Console of the Branch peer, select the **Configuration > Proxy Settings > MAPI Proxy** tab.
2. Select the **Enable acceleration for encrypted MAPI** option; the Domain alias list automatically populates with the alias created in "[Join the Branch Peer to the Primary Domain](#)" on page 285.
3. Click **Apply**.

Verify Encrypted MAPI Connections are Optimized

To verify that encrypted MAPI connections are being optimized:

- Initiate Outlook client-to-Exchange server actions, including emails with attachments. In the ProxySG appliance Management Console, monitor the Active Sessions (**Statistics > Sessions > Active Sessions**). The **Encrypted** label appends to connections intercepted and optimized by the ProxySG; for example: **MAPI 2007 (Encrypted)** shows in the **Details** column. In addition, the **P** (Protocol Optimization) column in Active Sessions should show a color (active)  icon.

A	S	FW	I	Duration	Client Bytes	Server Bytes	Savings	C	BC	OC	P	BM	Service Name	Protocol	Detail
				15 sec	816	240	70.59%						Endpoint Mapper	EPRI	Endpoint Mapper v3.0 Unknown Version
				15 sec	596	450	24.24%						Endpoint Mapper	EPRI	Endpoint Mapper v3.0 Unknown Version
				15 sec	2,820	2,966	0%						Endpoint Mapper	MSRPC	MS Exchange Directory NSPI Proxy v56.0 Unknown Version ...
				15 sec	2,704	2,842	0%						Endpoint Mapper	MSRPC	MS Exchange Directory NSPI Proxy v56.0 Unknown Version ...
				14 sec	66,609	36,990	46.24%						Endpoint Mapper	MAPI	MAPI 2007 (Encrypted)
				13 sec	3,748	3,946	0%						Endpoint Mapper	MSRPC	MS Exchange Directory NSPI Proxy v56.0 Unknown Version ...

If you misconfigure the deployment—for example, configure NTLM without Kerberos on the Exchange server—the ProxySG passes the connection through without optimization. If this occurs, the icon in the **P** column in Active Sessions is shown as inactive (gray). You should check the **Details** column for clues on why the connection wasn't optimized. For example, if the Domain Controller is offline or is unreachable by the Branch peer, the **Details** column displays "Unable to contact domain controller."

The following table lists the possible entries:

Active Session Detail Message	Reason
<i>Encrypted</i>	Encrypted MAPI connection is intercepted and optimized successfully
<i>Unable to contact domain controller</i>	The Domain Controller is offline or is unreachable by the Branch peer.
<i>Logon network security not set to negotiate on the client</i>	The Outlook account is not configured to use Negotiate Authentication (Outlook 2007 or newer) or Kerberos/NTLM Password Authentication (Outlook 2003 or older).
<i>Client security negotiation failed</i>	General error message.
<i>Server security negotiation failed</i>	General error message.
<i>Secure ADN not available</i>	- MAPI proxy failed to establish a secure ADN connection with the core ProxySG. - Outlook switched to a secure connection in the middle of conversation when the ADN tunnel was non secure.
<i>ADN tunnel is not encrypted</i>	Outlook switched to a secure connection in the middle of conversation when the ADN tunnel is not encrypted
<i>Encrypted MAPI not supported by peer SG</i>	Core ProxySG does not support encrypted MAPI protocol optimization

Active Session Detail Message	Reason
<i>NTLM-only client authentication type is unsupported</i>	The Outlook client has authenticated the connection with NTLM-only secure protocol. Protocol optimization is not supported.
<i>Kerberos-only client authentication type is unsupported</i>	The Outlook client has authenticated the connection with Kerberos-only secure protocol. Protocol optimization is not supported.
<i>Unexpected authentication type</i>	The Outlook client has authenticated the connection with an unexpected secure protocol. Protocol optimization is not supported.
<i>Unable to extract service principal name from SPNEGO connection</i>	Branch peer failed to extract exchangeMDB service principal name from SPNEGO packet which is required to negotiate Kerberos security context.
<i>Not intercepted by ADN concentrator</i>	If Branch peer is in standalone mode or failed to establish ADN connection with the Concentrator and Branch peers, the session downgrades to passthru mode.

- ❑ Display Errored Sessions (**Statistics > Sessions > Errored Sessions**) to investigate various MAPI issues related to client/server socket failures.

Section C: Intercept Skype for Business

The ProxySG will not be able to proxy some Skype for Business and Microsoft Lync application connections between clients once SSL Interception is enabled unless you follow all of the steps in this section. See the *Office 365 Best Practices* guide for additional information.

Skype for Business uses the following protocols (in addition to HTTPS):

- ❑ The Session Initiation Protocol (SIP) is commonly used for voice and video calls and instant messages. Because this protocol defines the messages and traffic between client endpoints, the ProxySG appliance interception of this traffic can cause dropped connections.
- ❑ The (Microsoft) Traversal Using Relay NAT (TURN) protocol is used to allocate a public IP address and port on a globally reachable server and relay media from one endpoint to another endpoint.

Configure ProxySG for for Skype and Lync Interception

Follow the instructions detailed in the *Office 365 Integration and Best Practices Webguide*, Skype for Business/Lync Fix section, to safely intercept Skype for Business and Microsoft Lync. Log in to [Symantec Product Documentation](#) to download the webguide.

Chapter 12: Managing the File Transport Protocol (FTP) Proxy

This chapter discusses the Blue Coat implementation of proxy support for File Transport Protocol (FTP).

Topics in this Chapter

This chapter includes information about the following topics:

- ["How Do I...?"](#)
- ["About FTP" on page 291](#)
- ["Configuring the ProxySG for Native FTP Proxy" on page 295](#)
- ["Configuring FTP Connection Welcome Banners" on page 298](#)
- ["Viewing FTP Statistics" on page 299](#)

How Do I...?

To use this chapter, identify the task and click the link:

How do I...	See...
Understand how the ProxySG manages IP addresses?	"About FTP" on page 291
Configure IP addresses?	"Configuring IP Addresses for FTP Control and Data Connections" on page 293
Configure native FTP?	"Configuring the ProxySG for Native FTP Proxy" on page 295
Configure Web FTP?	"About Web FTP" on page 156
Customize the <i>welcome banner</i> for the FTP proxy?	"Configuring FTP Connection Welcome Banners" on page 298
View FTP statistics?	"Viewing FTP Statistics" on page 299

About FTP

The ProxySG supports two FTP modes:

- **Native FTP**, where the client connects through the FTP proxy, either explicitly or transparently; the ProxySG then connects upstream through FTP (if necessary).

- ❑ **Web FTP**, where the client uses an explicit HTTP connection. Web FTP is used when a client connects in explicit mode using HTTP and accesses an `ftp://` URL. The ProxySG translates the HTTP request into an FTP request for the origin content server (OCS), if the content is not already cached, and then translates the FTP response with the file contents into an HTTP response for the client.

Native FTP uses two parallel TCP connections to transfer a file, a *control connection* and a *data connection*.

- ❑ **Control connections:** Used for sending commands and control information, such as user identification and password, between two hosts.
- ❑ **Data connections:** Used to send the file contents between two hosts. By default, the ProxySG allows both *active* and *passive* data connections.
 - **Active mode data connections:** Data connections initiated by an FTP server to an FTP client at the port and IP address requested by the FTP client. This type of connection method is useful when the FTP server can connect directly to the FTP client. The FTP command for active mode is PORT (for IPv4) or EPRT (for IPv6). When an IPv4 FTP client is communicating with an IPv6 FTP server, the ProxySG will perform the required conversion (PORT to EPRT); the clients and servers will be unaware that this conversion has taken place.
 - **Passive mode data connections:** Data connections initiated by an FTP client to an FTP server at the port and IP address requested by the FTP server. This type of connection is useful in situations where an FTP server is unable to make a direct connection to an FTP client because the client is located behind a firewall or other similar device where outbound connections from the client are allowed, but inbound connections to the client are blocked. The FTP command for passive mode is PASV (for IPv4) or EPSV (for IPv6). When an IPv4 FTP client is communicating with an IPv6 FTP server, the ProxySG will perform the required conversion (PASV to EPSV); the clients and servers will be unaware that this conversion has taken place.

Note: When using the FTP in active mode, the FTP data connection is formed from the server (OCS) to the client, which is opposite from the direction of the FTP control connection. As a result, when the FTP connections are enabled for ADN, the roles of the Branch and Concentrator for the data connection are in reverse of those used for the control connection. The type of ADN tunnel (Explicit, Translucent or Transparent) set up for the data connection is therefore dictated by the tunnel mode configuration, which can be used for any connection from the server to the client that needs to go over ADN. For more information, see "[Configuring the Tunnel Mode](#)" on page 730.

For example, if the control connection for an Active mode FTP uses explicit ADN tunnels, it is possible that the data connection that goes from the server

to the client is transparent. To use explicit connections for the FTP data connection as well, it might be necessary to advertise the FTP client's subnet address on the ProxySG appliance intercepting the FTP connection.

This section discusses:

- ❑ "Configuring IP Addresses for FTP Control and Data Connections" on page 293
- ❑ "Client-Side Data Connections Mode" on page 294
- ❑ "FTP Server Notes" on page 295

Configuring IP Addresses for FTP Control and Data Connections

The FTP client determines whether the client-side data connection is active or passive from the client to the ProxySG. The ProxySG determines the server-side connections.

By default, the ProxySG allows both active and passive data mode connections. FTP connections are divided into client-side control and data connections and server-side control and data connections.

- ❑ **Client-side control connection:** The proxy always uses the client's IP address to respond to the client. No configuration is necessary here.
- ❑ **Client-side data connection:** The proxy's behavior depends on the `ftp.match_client_data_ip(yes | no)` property that is set via policy using CPL. If this property is enabled (the default), the proxy uses the same IP address for the data connection as it uses for the client-side control connection. If the property is disabled, the proxy uses its own IP address, choosing the address associated with the interface used to connect back to the client.

When an FTP client uses different protocols for control and data connections (for example, IPv4 for control and IPv6 for data), the

`ftp.match_client_data_ip` property must be set to **no** so that the ProxySG's address is used for the data connection. Because each ProxySG interface is configured with an IPv4 and an IPv6 address in a mixed Internet protocol environment, the ProxySG will use the appropriate IP address for the type of FTP server. For example, for transferring data to an IPv6 FTP server, the ProxySG will set up with the data connection using its IPv6 address.

When the client-side data and control connections are over IPv4 and the server-side control and data connections are over IPv6, the

`ftp.match_client_data_ip` property can be set to **yes**.

- ❑ **Server-side control connection:** The proxy uses the IP address selected by the `reflect_ip(auto | no | client | vip | ip_address)` property. By default, this is the local proxy IP address associated with the interface used to connect to the server.

Client IP reflection is set globally from the **Configuration > Proxy Settings > General** tab. By default, the CPL `reflect_ip( )` setting is **auto**, which uses this global configuration value.

Client IP reflection will automatically be disabled when the client is IPv4 and the server is IPv6.

Note: Setting client IP address reflection for FTP affects the source address that is used when making the outgoing control connection to the origin server. It might also affect which address is used by the proxy for data connections.

- ❑ **Server-side data connection:** The proxy's behavior depends on the `ftp.match_server_data_ip(yes | no)` property. If this property is enabled (the default), the proxy uses the same IP address for the data connection as it used for the server-side control connection. If the property is disabled, the proxy uses its own IP address to communicate with the server, choosing the address associated with the interface used to connect to the server.

Note: Either the `reflect_ip( )` property or the `reflect-client-ip` configuration must be set for the `ftp.match_server_data_ip(yes)` property to be meaningful.

For information on creating and modifying policy through VPM, refer to the *Visual Policy Manager Reference*. For information on creating and modifying policy through CPL, refer to the *Content Policy Language Reference*. The `ftp.match_server_data_ip( )` and `ftp.match_client_data_ip( )` properties can only be set through CPL.

Client-Side Data Connections Mode

Administrators determine how the ProxySG responds to a request from an FTP client for a passive mode data connection.

By default, some FTP clients do not open a passive mode data connection to an IP address that is different from the IP address used for the control connection.

When passive mode is disabled, some FTP clients try a PORT (IPv4) or EPRT (IPv6) command automatically, which allows requests to be received when the client doesn't allow passive connections to a different IP address.

Note: Some clients might display an error when passive mode is disabled on the ProxySG, requiring you to manually request active mode using the PORT/EPRT FTP commands.

The FTP client software controls any messages displayed to the end user as a result of this response from the ProxySG.

Server-Side Data Connections Mode

The `ftp.server_data(auto | passive | port)` property controls the type of server-side data connection that the ProxySG opens to the server. The default of `auto` means to try a passive connection first and then fall back to an active connection if that fails.

FTP Server Notes

IIS and WS_FTP servers do not support:

- ❑ Passive data connections with a source IP address that is different from the source IP address of the control connection.
- ❑ Active data connections with a destination IP address that differs from the source IP address of the control connection.

The `ftp.match_server_data_ip(no)` property most likely will not work correctly with these servers.

Notes

- ❑ Internet Explorer does not support proxy authentication for native FTP.
- ❑ The FTP proxy does not support customized exception text; that is, you can use policy to deny requests, but you can't control the text sent in the error message.

Configuring the ProxySG for Native FTP Proxy

This section discusses:

- ❑ ["Modifying the FTP Proxy Service"](#)
- ❑ ["Configuring the FTP Proxy"](#) on page 296
- ❑ ["Configuring FTP Clients for Explicit Proxy"](#) on page 297


Modifying the FTP Proxy Service

To use the capabilities of the FTP proxy, you need to make sure the FTP service is set to intercept traffic. The following procedure describes how to verify this setting, and explains other attributes within the service.

Note: Web FTP requires an HTTP service, not an FTP service. For information on configuring an HTTP proxy service, see [Chapter 8: "Intercepting and Optimizing HTTP Traffic"](#) on page 153.

To modify the FTP proxy service:

1. From the Management Console, select **Configuration > Services > Proxy Services**.



2. Intercept FTP traffic:

- Scroll the list of service groups and click **Standard** to expand the services list and locate the FTP service.
- Is the FTP service set to **Bypass** or **Intercept**? If necessary, select **Intercept** from the drop-down list.

3. Click **Apply**.


Now that you have verified that the ProxySG is intercepting FTP traffic, configure the FTP proxy options. Proceed to "[Configuring the FTP Proxy](#)" on page 296.

Configuring the FTP Proxy

The FTP proxy has several configurable settings related to caching of FTP objects and whether passive mode is allowed.

To configure the FTP proxy:

- Select **Configuration > Proxy Settings > FTP Proxy**.



2. Configure the FTP proxy settings:

- a. Select **Allow caching of FTP objects**. The default is enabled.
- b. Determine how long the object will be cached, in relation to when it was last modified. This setting assumes the object's last-modified date/time is available from the server. (The next setting, in step c below, applies to situations when the last-modified date is unknown.) The default is 10%. The amount of time that the object will be cached is calculated as follows:

$$\text{percentage} * (\text{current\_time} - \text{last\_modified\_time})$$

where `current_time` is the time when the object was requested by the client. So, if it's been 10 days since the object was modified, and the setting is 10%, the object will be cached for one day.

- c. Enter an amount, in hours, that the object remains in the cache before becoming eligible for deletion. This setting applies to objects for which the last-modified date is unknown. The default is 24 hours.
- d. Select **Allow use of passive mode to clients**. The default is **enabled**, allowing data connections to be initiated by an FTP client to an FTP server at the port and IP address requested by the FTP server. (Active mode connections are always allowed, regardless of whether the passive mode setting is enabled or disabled.)
- e. (Optional) See ["Configuring FTP Connection Welcome Banners"](#) on page 298.

3. Click **Apply**.

Note: Neither proxy authentication for transparent FTP nor proxy chaining are supported with the Checkpoint syntax. When native FTP traffic from an FTP client (such as WSFtp) is being authenticated by the ProxySG using the Raptor syntax, the recommended authentication mode is auto or proxy.

Configuring FTP Clients for Explicit Proxy

To explicitly proxy to the ProxySG, each FTP client must be configured with the IP address of the ProxySG. In addition, the client may need additional configuration. The example below describes how to configure the WSFtp client; you will want to use equivalent steps for other FTP clients.

- ☐ Enable firewall.

- ❑ Select **USER with no logon** unless you are doing proxy authentication. In that case, select **USER remotelD@remoteHost firelD** and specify a proxy username and password.




Figure 12-1 Example WSFTtp Client Configuration


Configuring FTP Connection Welcome Banners

You can customize banners that usually describe the policies and content of the FTP server displayed to FTP clients. Without modification, the ProxySG sends a default banner to newly-connected FTP clients: **Welcome to Blue Coat FTP**. However, you might not want users to know that a ProxySG exists on the network. A default banner can be defined in the Management Console or the CLI, but other banners defined for specific groups can be created in policy layers.

Note: Configurable banners are only displayed when FTP is explicitly proxied through the ProxySG. In transparent deployments, the banner is sent to the client when proxy authentication is required; otherwise, the FTP server sends the banner.

To define the default FTP banner:

1. Select **Configuration > Services > FTP Proxy**.
2. In the **Welcome Banner** field, enter a line of text that is displayed on FTP clients upon connection. If the message length spans multiple lines, the ProxySG automatically formats the string for multiline capability.



The welcome banner text is overridden by the policy property `ftp.welcome_banner()`. This is required for explicit proxy requests, when doing proxy authentication, and also when the policy property `ftp.server_connection(deferred|immediate)` is set to defer the connection.

3. Click **Apply**.

Viewing FTP Statistics

See "[HTTP/FTP History Statistics](#)" on page 202 for information about viewing the FTP statistics.

Chapter 13: Managing the Domain Name Service (DNS) Proxy

This chapter discusses managing Domain Name Service (DNS) traffic through the DNS proxy on the ProxySG (to configure the ProxySG connections to DNS servers, see ["Adding DNS Servers to the Primary or Alternate Group"](#) on page 816).

Topics in this Chapter

This chapter includes information about the following topics:

- ["About the DNS Proxy"](#)
- ["Intercepting the DNS Proxy Service"](#) on page 302
- ["Creating a Resolving Name List"](#) on page 302

About the DNS Proxy

The ProxySG is *not* a DNS server. It does not perform zone transfers, and it forwards recursive queries to other name servers. When a DNS proxy *service* is enabled (*intercepted*), it listens on port 53 for both explicit and transparent DNS domain query requests. By default, the service is created but not enabled.

The DNS proxy performs a lookup of the DNS cache on the ProxySG to determine if requests can be answered locally. If yes, the ProxySG responds to the DNS request. If not, the DNS proxy forwards the request to the DNS server list configured on the ProxySG.

Through policy, you can configure a list of resolved domain names (the *resolving name list*) the DNS uses. The domain name in each query received by the ProxySG is compared against the resolving name list. Upon a match, the ProxySG checks the resolving list. If a domain name match is found but no IP address was configured for the domain, the ProxySG sends a DNS query response containing its own IP address. If a domain name match is found with a corresponding IP address, that IP address is returned in a DNS query response. All unmatched queries are sent to the name servers configured on the ProxySG.

IPv6 Support

The DNS proxy is able to communicate using IPv4 or IPv6, either explicitly or transparently.

The resolving name list can contain entries for IPv4 and IPv6 addresses. An entry can contain either IPv4 or IPv6 addresses, although you cannot combine IPv4 and IPv6 addresses in a single entry.


Intercepting the DNS Proxy Service

By default (upon upgrade and on new systems), the ProxySG has an DNS proxy service configured on port 53. The service is configured to listen to all IP addresses, but is set to Bypass mode.

The following procedure describes how to change the service to Intercept mode.

To configure the DNS proxy to intercept traffic:

1. From the Management Console, select **Configuration > Services > Proxy Services**.



2. Intercept DNS traffic:
 - a. Expand the **Standard** service group, and select **DNS**.
 - b. If the DNS service is currently set to **Bypass**, select **Intercept**.
3. Click **Apply**.

Creating a Resolving Name List

You can create the resolving name list that the DNS proxy uses to resolve domain names. This procedure can only be done through policy. (For a discussion on using the <DNS-Proxy> layer, refer to the <Emphasis>Blue Coat Content Policy Language Guide.</Emphasis>)

Each entry in the list contains a domain-name matching pattern. The matching rules are:

- ❑ `test.com` matches only `test.com` and nothing else.
- ❑ `.test.com` matches `test.com`, `www.test.com` and so on.
- ❑ `“.”` matches all domain names.

An optional IP address can be added, which allows the DNS proxy to return any IP address if the DNS request's name matches the domain name suffix string (`domain.name`). Either IPv4 or IPv6 addresses can be specified.

To create a resolving name list, create a policy, using the <DNS-Proxy> layer, that contains text similar to the following:

```
<DNS-Proxy>
  dns.request.name=www.example.com dns.respond.a(vip)
-or-
<DNS-Proxy>
  dns.request.name=.example.com dns.respond.a(vip)
-or-
<DNS-Proxy>
  dns.request.name=www.example.com dns.respond.a(10.1.2.3)
-or-
<DNS-Proxy>
  dns.request.name=www.google.com dns.respond.aaaa(2001::1)
```

An entry can contain either IPv4 or IPv6 addresses, although you cannot combine IPv4 and IPv6 addresses in a single entry. Use the `dns.respond.a` property for IPv4 hosts and `dns.respond.aaaa` for IPv6 hosts. If you specify `vip` instead of a specific IP address, the response will contain the ProxySG IP address (the IPv6 address for `dns.respond.aaaa` or the IPv4 address for `dns.respond.a`).

Note: You can also create a resolving name list using the Visual Policy Manager (VPM). For more information about the **DNS Access Layer** in the VPM, refer to the *Visual Policy Manager Reference*.

Chapter 14: Managing a SOCKS Proxy

This chapter discusses the ProxySG SOCKS proxy.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ ["About SOCKS Deployments"](#) on page 305
- ❑ ["Intercepting the SOCKS Proxy Service"](#) on page 306
- ❑ ["Configuring the SOCKS Proxy"](#) on page 307
- ❑ ["Using Policy to Control the SOCKS Proxy"](#) on page 307
- ❑ ["Viewing SOCKS History Statistics"](#) on page 308
- ❑ ["Viewing SOCKS History Statistics"](#) on page 308

About SOCKS Deployments

While SOCKS servers are generally used to provide firewall protection to an enterprise, they also can be used to provide a generic way to proxy any TCP/IP or UDP protocols. The ProxySG supports both SOCKS v4/4a and SOCKS v5; however, because of increased username and password authentication capabilities and compression support, Blue Coat recommends that you use SOCKS v5. Note that there is only one listener for all SOCKS connections (SOCKS v4 and v5).

In a typical MACH5 deployment, the SOCKS proxy works with the Endpoint Mapper proxy and MAPI handoff. In this deployment, you will:

- ❑ Create an Endpoint Mapper proxy at the remote office (the downstream proxy) that intercepts Microsoft RPC traffic and creates dynamic TCP tunnels. Traffic to port 135 is transparently redirected to this service using bridging or L4 switch or WCCP. For information on creating and enabling an Endpoint Mapper proxy service, see [Chapter 11: "Managing Outlook365 Applications"](#) on page 271.
- ❑ Create any other TCP tunnel proxies you need at the remote office: SMTP, DNS, and the like. For information on configuring TCP tunnels, see [Section C: "Creating Custom Proxy Services"](#) on page 120.
- ❑ Create a SOCKS gateway at the remote office and enable compression for that gateway. This SOCKS gateway points to a SOCKS proxy located at the main office location (the upstream proxy, the core of the network). For information on creating a SOCKS gateway and enabling SOCKS compression, see [Chapter 41: "SOCKS Gateway Configuration"](#) on page 839.

- ❑ Set policy to forward TCP traffic through that SOCKS gateway. You can do this through the <proxy> layer using either the VPM or CPL. For more information, see "Using Policy to Control the SOCKS Proxy" on page 307.

IPv6 Support


The SOCKS proxy includes basic IPv6 support for CONNECT and BIND.

In addition, for any service that uses the SOCKS proxy, you can create listeners that bypass or intercept connections for IPv6 sources or destinations.

Intercepting the SOCKS Proxy Service

By default, the ProxySG includes a pre-defined service for SOCKS that listens on port 1080, but the service is set to **Bypass**. To use the SOCKS proxy, you need to set this service to **Intercept**. To configure the SOCKS proxy to intercept traffic:

1. In the Management Console, select **Configuration > Services > Proxy Services**.



2. Change the SOCKS service to intercept
 - a. Scroll the list of service groups, click **Standard**, and select **SOCKS**.
 - b. If the **Action** for the default service (port 1080) is set to **Bypass**, select **Intercept** from the drop-down list.
3. Click **Apply**.

Configuring the SOCKS Proxy

Complete the following steps to create a SOCKS proxy and to configure SOCKS-proxy connection and timeout values. For more information, see "[About SOCKS Deployments](#)" on page 305.

To create a SOCKS proxy server:

1. Select **Configuration > Services > SOCKS Proxy**.

The screenshot shows the 'SOCKS Proxy' configuration window. Under the 'SOCKS proxy options' section, the following values are displayed:

Option	Value	Description
Max-Connections:	0	(0 means unlimited)
Connection timeout:	120	seconds
Bind timeout on accept:	120	seconds
Minimum idle timeout:	7200	seconds (0 means unlimited)
Maximum idle timeout:	0	seconds (0 means unlimited)

2. The displayed defaults should be sufficient for most purposes. The following table discusses the options.

Table 14–1 SOCKS Proxy Options

Option	Suboption	Description
Max-Connections	<i>connections</i>	Set maximum allowed SOCKS client connections. The default of 0 indicates an infinite number of connections are allowed.
Connection timeout	<i>seconds</i>	Set maximum time to wait on an outbound CONNECT.
Bind timeout on accept	<i>seconds</i>	Set maximum time to wait on an inbound BIND.
Minimum idle timeout	<i>seconds</i>	Specifies the minimum timeout after which SOCKS can consider the connection for termination when the maximum connections are reached.
Maximum idle timeout	<i>seconds</i>	Specifies the max idle timeout value after which SOCKS terminates the connection.

Using Policy to Control the SOCKS Proxy

After the basic configuration for the SOCKS proxy has been set, you can define policy to control the SOCKS proxy.

- ❑ To use SOCKS version 5, which allows you to use a SOCKS username/password, you must set the version through policy.
 - If using the VPM, go to a **Forwarding Layer**, select **Source > Set Source Object > New > SOCKS Version**.
 - If using CPL, enter the following:

```
<proxy> client.protocol=socks
ALLOW socks.version=5
DENY
```

- ❑ If browsers and FTP clients are configured to use SOCKS encapsulation and a rule in policy is matched that denies a transaction, a **page cannot be displayed** message displays instead of an exception page.

This is expected behavior, as a deny action abruptly closes the client's TCP connection, yet the client is expecting a SOCKS-style closure of the connection. You can avoid this, and return an exception page, by applying the following policy:

- If using the VPM, go to a **Web Access Layer**, create two rules. For the first rule, select **Service > New > Client Protocol > SOCKS > TCP Tunneling over SOCKS**; for the second, select **Service > New > Client Protocol > SOCKS > All SOCKS**.
- If using CPL, enter the following:

```
<Proxy>
DENY socks=yes tunneled=yes
DENY socks=yes
```

Viewing SOCKS History Statistics

The **SOCKS History** tabs (SOCKS Clients, SOCKS Connections, and SOCKS client and server compression) display client data, Connect, Bind, and UPD Associate requests, client and server UDP, TCP and compression requests.


Note: The SOCKS history statistics are available only through the Management Console.

Viewing SOCKS Clients

The SOCKS Clients tab displays SOCKS Client data.

To view SOCKS client data:

1. Select **Statistics > SOCKS History > SOCKS Clients**.
2. Select a time period for the graph from the **Duration:** drop-down list.
3. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.




Viewing SOCKS Connections

The SOCKS Connections tab displays SOCKS Connection data.

To view SOCKS connection data:

Select **Statistics > SOCKS History > SOCKS Connections**.



Viewing SOCKS Client and Server Compression Gain Statistics

You can view SOCKS client and server compression-gain statistics for the ProxySG over the last 60 minutes, 24 hours, and 30 days in the **Client Comp. Gain** and the **Server Comp. Gain** tabs. These statistics are not available through the CLI.

The green display on the bar graph represents uncompressed data; the blue display represents compressed data. Hover your cursor over the graph to see the compressed gain data.


See one of the following topics:

- ❑ "Viewing SOCKS Client Compressed Gain Statistics"
- ❑ "Viewing SOCKS Server Compressed Gain Statistics"

Viewing SOCKS Client Compressed Gain Statistics

To view SOCKS client compressed gain statistics:

1. Select **Statistics > SOCKS History > Client Comp. Gain**.
2. Select a time period for the graph from the **Duration:** drop-down list.




3. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.

Viewing SOCKS Server Compressed Gain Statistics

To view SOCKS Server compressed gain statistics:

1. Select **Statistics > SOCKS History > Server Comp. Gain**.
2. Select a time period from the **Duration:** drop-down list.



3. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.

Chapter 15: Managing Shell Proxies

This chapter discusses how to configure the Telnet shell proxy. Shell proxies provide shells which allow a client to connect to the ProxySG. In this version, only a Telnet shell proxy is supported.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ ["About Shell Proxies" on page 313](#)
- ❑ ["Customizing Policy Settings for Shell Proxies" on page 314](#)
- ❑ ["About Telnet Shell Proxies" on page 314](#)
- ❑ ["Configuring the Telnet Shell Proxy Service Options" on page 315](#)
- ❑ ["Viewing Shell History Statistics" on page 317](#)

About Shell Proxies

Using a shell proxy, you can:

- ❑ terminate a Telnet protocol connection either transparently or explicitly.
- ❑ authenticate users either transparently or explicitly.
- ❑ view the access log.
- ❑ enforce policies specified by CPL.
- ❑ communicate through an upstream SOCKS gateway and HTTP proxy using the `CONNECT` method.

Within the shell, you can configure the prompt and various banners using CPL `$substitutions`. You can also use hard-coded text instead of CPL substitutions (available substitutions are listed in the table below). The syntax for a CPL substitution is:

```
$(CPL_property)
```

Table 15–1 CPL Substitutions for Shell Proxies

Substitution	Description
<code>proxy.name</code> or <code>appliance.name</code>	Configured name of the ProxySG.
<code>proxy.address</code>	IP address of the appliance on which this connection is accepted. You can specify either an IPv4 or an IPv6 address.
<code>proxy.card</code>	Adapter number of the appliance on which this connection is accepted.

Table 15–1 CPL Substitutions for Shell Proxies

Substitution	Description
<code>client.protocol</code>	This is <code>telnet</code> .
<code>client.address</code>	IP address of the client. IPv4 and IPv6 addresses are accepted.
<code>proxy.primary_address</code> or <code>appliance.primary_address</code>	Primary address of the proxy, not where the user is connected. You can specify either an IPv4 or an IPv6 address.
<code>release.id</code>	SGOS version.

Customizing Policy Settings for Shell Proxies

For information on using CPL to manage shell proxies, refer to <Emphasis>Blue Coat Content Policy Language Guide.

Boundary Conditions for Shell Proxies

- ❑ A hardcoded timeout of five minutes is enforced from the acceptance of a new connection until destination information is provided using the Telnet command.
- ❑ If proxy authentication is enabled, users have three chances to provide correct credentials.
- ❑ Users are not authenticated until destination information is provided.
- ❑ Users can only enter up to an accumulated 2048 characters while providing the destination information. (Previous attempts count against the total number of characters.)
- ❑ Connection to an upstream HTTP proxy is not encouraged.
- ❑ If connections from untrustworthy IP address or subnet are not desired, then a client IP/subnet-based *deny* policy must be written.

About Telnet Shell Proxies

The Telnet shell proxy allows you to manage a Telnet protocol connection to the ProxySG. Using the Telnet shell proxy, the ProxySG performs:

- ❑ Explicit termination without proxy authentication, where you explicitly connect through Telnet to the ProxySG host name or IP address. In this case, the ProxySG provides a shell.
- ❑ Explicit termination with proxy authentication, where after obtaining the destination host and port information from the user, the ProxySG challenges for proxy credentials. After the correct proxy credentials are provided and authenticated, the appliance makes an upstream connection and goes into tunnel mode. In this case, the appliance provides a shell.

- ❑ Transparent termination without proxy authentication, where the ProxySG intercepts Telnet traffic through an L4 switch, software bridge, or any other transparent redirection mechanism. From the destination address of TCP socket, the ProxySG obtains OCS contact information and makes the appropriate upstream connection, either directly or through any configured proxy. For more information on configuring a transparent proxy, see [Chapter 6: "Explicit and Transparent Proxy"](#) on page 99.
- ❑ Transparent termination with proxy authentication, where, after intercepting the transparent connection, the ProxySG challenges for proxy credentials. After the correct proxy credentials are provided and authenticated, the ProxySG makes an upstream connection and goes into tunnel mode.

After in the shell, the following commands are available:

- ❑ `help`: Displays available commands and their effects.
- ❑ `telnet server[:port]`: Makes an outgoing Telnet connection to the specified server. The colon (:) between server and port can be replaced with a space, if preferred. The server can be an IPv4 or an IPv6 host.
- ❑ `exit`: Terminates the shell session.

Configuring the Telnet Shell Proxy Service Options

This section describes how to change the default service options and add new services.


Changing the Telnet Shell Proxy Service to Intercept All IP Addresses on Port 23

The service is configured to listen to all IP addresses, but is set in Bypass mode. The following procedure describes how to change the service to Intercept mode. Default settings are:

- ❑ Proxy Edition—a Telnet proxy service is configured but disabled on port 23 on a new system.
- ❑ Proxy Edition— a Telnet proxy service is not created on an upgrade.
- ❑ MACH5 Edition—a transparent TCP tunnel connection listening on port 23 is created in place of the default Telnet proxy service.

To configure the Telnet Shell proxy to intercept traffic:

1. From the Management Console, select **Configuration > Services > Proxy Services**.




2. Scroll to the **Bypass Recommended** service group and click it to expand the list.
3. Select **Telnet**.
4. From the drop-down list, select **Intercept**.
5. Click **Apply**.

Customizing Welcome and Realm Banners and Prompt Settings

You can configure banners for the Telnet shell and the realm and set the prompt that users see when entering the shell.

To customize Telnet shell proxy settings:

1. Select **Configuration > Proxy Settings > Shell Proxies > Telnet Proxy Settings**.



2. To set the maximum concurrent connections, select **Limit Max Connections**. Enter the number of maximum concurrent connections allowed for this service. Allowed values are between 1 and 65535.

3. (Optional) Change the default banner settings.

- **Welcome banner**—Users see this when they enter the shell. The default string is: `Blue Coat $(module_name) proxy.`
- **Realm banner**—Users see this help message just before they see the **Username** prompt for proxy authentication. The default string is:
`Enter credentials for realm $(realm).`
- **Prompt**—The command prompt. The default string is:
`$(module_name)-proxy>.`

For a list of available substitutions, see "[Customizing Policy Settings for Shell Proxies](#)" on page 314.

Click **View/Edit** to display the respective banner dialog. Change the string. Click **OK**.

4. Click **Apply**.**Notes for Telnet Shell Proxies**

- Telnet credential exchange is in plaintext.
- A Telnet proxy cannot be used to communicate with non-Telnet servers (such as Web servers on port 80) because Telnet proxies negotiate Telnet options with the client before a server connection can be established.


Viewing Shell History Statistics

The **Shell History** tab displays client connections over the last 60-minute, 24-hour, and 30-day period.

Note: The Shell history statistics are available only through the Management Console.

To view Shell history statistics:

1. Select **Statistics > Protocol Details > Shell History**.



2. Select a time- period for the graph from the **Duration:** drop-down list. The default setting is last hour.
3. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.

Chapter 16: Configuring and Managing an HTTPS Reverse Proxy

This section describes how to use the Blue Coat HTTPS Reverse Proxy solution. It includes the following topics:

- ❑ [Section A: "About the HTTPS Reverse Proxy" on page 319](#)
- ❑ [Section B: "Configuring the HTTPS Reverse Proxy" on page 320](#)
- ❑ [Section C: "Configuring HTTP or HTTPS Origination to the Origin Content Server" on page 327](#)

Section A: About the HTTPS Reverse Proxy

The Blue Coat HTTPS Reverse Proxy implementation:

- ❑ Combines hardware-based SSL acceleration with full caching functionality.
- ❑ Establishes and services incoming SSL sessions.
- ❑ Provides TLS v1.2, TLS v1.1, TLSv1, SSL v3.0, SSL v2.0 and protocol support.
- ❑ Supports IPv6 connections.

Section B: Configuring the HTTPS Reverse Proxy

This section describes how to enable the HTTPS reverse proxy.

Note: One common scenario in using HTTPS reverse proxy, which connects the client to the ProxySG, is in conjunction with HTTPS *origination*, which is used to connect to the origin content server (OCS). For more information on this option, see [Section C: "Configuring HTTP or HTTPS Origination to the Origin Content Server"](#) on page 327.

Prerequisite Tasks

Before creating an HTTP reverse proxy service, you must:

- ❑ Create or import a keyring (**Configuration > SSL > Keyrings > SSL Keyrings**).
- ❑ (If necessary) Create a Certificate Signing Request (CSR) that can be sent to a Certificate Signing Authority (CA). After the CSR has been signed and the certificate has been created, import the certificate to the keyring you created or imported in the previous step. (Select **Configuration > SSL > Keyrings**. Select the keyring you created or imported, and then click **Edit**. In the Certificate Signing Request section, click **Import**, paste the CSR, and click **OK**. Click **Close > Apply**).

-or-

- ❑ Create a certificate for internal use and associate it with the keyring.
- ❑ (Optional, if using server certificates from CAs) Import Certificate Revocation Lists (CRLs) so the ProxySG can verify that certificates are still valid.


When these steps are complete, you can configure the HTTPS reverse proxy service.

Creating an HTTPS Reverse Proxy Service

Unlike other services, the ProxySG does not create an HTTPS Reverse Proxy service by default. (The ProxySG has an HTTPS proxy service configured on port 443.) Therefore, you must create a new service.

To create an HTTPS reverse proxy service:

1. Select **Configuration > Services > Proxy Services**.



2. Click **New Service**.

The screenshot shows the 'New Service' configuration window. It includes the following fields and sections:

- Name:** RevProxy
- Service Group:** Standard
- Proxy settings:**
 - Proxy:** HTTPS Reverse Proxy
 - Keyring:** default
 - CCL:** <All CA Certificates>
 - SSL protocols:**
 - ☒ TLSv1.2
 - ☒ TLSv1.1
 - ☒ TLSv1
 - ☐ SSLv3*
 - ☐ SSLv2*
 - * These SSL protocols are not recommended.
 - ☒ Verify Client
 - ☒ Forward Client Cert
- TCP/IP Settings:**
 - ☒ Early Intercept
- Application Delivery Network Settings:**
 - ☒ Enable ADN
 - ☒ Enable byte caching
 - ☒ Enable compression
 - Retention priority:** normal
- Listeners:**


Source IP	Destination IP	Port range	Action
- Buttons:** New, Edit, Delete

3. In the **Name** field, enter a name to identify the service.
4. From the **Service Group** drop-down list, select which group displays the service on the main page. You can add the service to a default group or any already created custom groups.
5. Configure **Proxy Settings** options:
 - a. Select **HTTPS Reverse Proxy** from the **Proxy settings** drop-down list.
 - b. In the **Keyring** drop-down list, select any already created keyring that is on the system. The system ships with a default keyring that is reusable for each HTTPS service.

Note: The **configuration-passwords-key** keyring that shipped with the ProxySG does *not* contain a certificate.

The **appliance-key** keyring does contain a certificate if you have Internet connectivity, but it cannot be used for purposes other than appliance authentication. For information about appliance authentication, see the Authentication topics.

- c. **CA Cert List:** Use the drop-down list to select any already created list that is on the system.
 - d. **SSL Versions:** Select the version(s) to use for this service from the list. The default is **TLS v1**, **TLS v1.1**, and **TLS v1.2**.
 - e. **Verify Client:** Select this option to enable mutual SSL authentication. See ["About Mutual SSL Authentication"](#) on page 325 for information.
Selecting this option makes the **Forward Client Certificate** option available.
 - f. **Forward Client Cert:** (Available if **Verify Client** is selected) Select this option to put the extracted client certificate information into the `Client-Cert` header that is included in the request when it is forwarded to the origin content server. The header contains the certificate serial number, subject, validity dates, and issuer (all as `name=value` pairs). The actual certificate itself is not forwarded.
6. Configure **Application Delivery Network** options:
- a. **Enable ADN:** Enabling ADN does not guarantee acceleration—the actual enable decision is determined by ADN routing (for explicit deployment) and network setup (for transparent deployment)
 - b. The **Compression and Byte Caching** options are selected by default if you enabled ADN optimization during initial configuration. Clear these options if you are not configuring ADN optimization.
7. Create a listener for the IP address(es) and ports that this application protocol uses. In the **Listeners** area, click **New**. The **New Listener** dialog displays.



8. Configure the new listener attributes:
 - a. In the **Source address** area, the most common selection is **All**, which means the service applies to requests from any client (IPv4 or IPv6). You can, however, restrict this listener to a specific IPv4/IPv6 address or user subnet/prefix length.
 - b. Select a **Destination address** from the options. The correct selection might depend on network configuration. For descriptions of the options, see ["About Proxy Services"](#) on page 110.
 - c. In the **Port Range** field, the default port is 443. Only change this if your network uses another port for SSL.
 - d. In the **Action** area, select **Intercept**.
 - e. Click **OK** to close the dialog.
9. Click **Ok** to add the new service to the selected service group.
10. Click **Apply**.

About Mutual SSL Authentication

During an SSL handshake, the client and server negotiate the mode of operation, the type of authentication required by both parties, the cryptographic and hashing algorithms to use for providing confidentiality and integrity, and the compression algorithm to use for the session.

SSL authentication can use the following modes of operation:


- ❑ Typical SSL authentication: This mode provides confidentiality and integrity of the data sent between the client and the server, and requires the server to authenticate to the client using an X.509 certificate.
- ❑ Mutual SSL authentication: In this mode, the server authenticates to the client using an X.509 certificate and the client must authenticate to the server with a separate X.509 certificate.

When a Common Access Card (CAC) is used, the certificate identifies the user who owns the CAC. For information on CAC authentication, refer to the *Common Access Card Solutions Guide*.

In mutual SSL authentication, an SSL connection between a client and a server is established only if the client and server validate each other's identity during the SSL handshake. Both the server and the client must have their own valid certificate and the associated private key in order to authenticate.

Note: TLS is supported based on the server and client in use. For brevity, this section refers only to SSL; however, SSL can be used interchangeably with TLS.

Typical SSL Authentication




In this scenario, the user logs in to the ProxySG appliance (server) using a browser (client). During this process, the client (browser) validates the server (ProxySG appliance) certificate. This includes the following checks:

- ❑ The certificate subject must match the server's hostname.
- ❑ The certificate must be issued by a CA listed in the browser's Trusted Root Certificate store.

- The client confirms that the server has the certificate's private key by challenging the server to sign random data. The client validates the signature using the server's certificate.

Mutual SSL Authentication



In this scenario, the user logs in to the ProxySG appliance using mutual SSL authentication. During this process:

1. The client (browser) validates the server (ProxySG appliance) certificate. This includes the following checks:
 - The certificate subject must match the server's hostname.
 - The certificate must be issued by a CA listed in the browser's Trusted Root Certificate store.
 - The client confirms that the server has the certificate's private key by challenging the server to sign random data. The client validates the signature using the server's certificate.
2. The server (ProxySG appliance) validates the client certificate that the browser presents. This includes the following checks:
 - The certificate must be issued by a CA in the CCL for the ProxySG appliance service that is performing the validation.
 - The server confirms that the client has the certificate's private key by challenging the client to sign random data. The server validates the signature using the client's certificate.
 - The certificate must be valid; it must have a valid signature and not be expired.
 - (If using a CRL) The certificate must not have been revoked.

Section C: Configuring HTTP or HTTPS Origination to the Origin Content Server

In previous procedures, you configured HTTPS Reverse Proxy to the ProxySG. In two common termination scenarios, you must also configure HTTPS origination to the Origin Content Server (OCS).

The first two scenarios are used to provide a secure connection between the proxy and server, if, for example, the proxy is in a branch office and is not co-located with the server.

Table 16–1 Scenario 1: HTTPS Reverse Proxy with HTTPS Origination

HTTPS Reverse Proxy	HTTPS Origination
Client > HTTPS > ProxySG	ProxySG > HTTPS > Origin Content Server
Steps • Configure a keyring. • Configure the SSL client. • Configure the HTTPS service.	Steps • (Optional) Add a forwarding host. • (Optional) Set an HTTPS port. • (Optional) Enable server certificate

Table 16–2 Scenario 2: HTTP Termination with HTTPS Origination

HTTPS Reverse Proxy	HTTPS Origination
Client > HTTPS > ProxySG	ProxySG > HTTPS > Origin Content Server
Steps • Client is explicitly proxied.	Steps • Server URL rewrite. -or- • Add a forwarding host • Set an HTTPS port. • (Optional) Enable server certificate verification

Using server URL rewrite is the preferred method. For information on rewriting the server URL, refer to the *<Emphasis>Blue Coat Content Policy Language Guide*.

To configure HTTPS origination:

At the (config) command prompt, enter the following commands:

```
SGOS#(config forwarding) create host_alias hostname
https[=port_number] server ssl-verify-server=yes
```

where:

Table 16–3 HTTPS Origination Commands

Option	Parameters	Description
<i>host_alias</i>	<i>alias_name</i>	Specifies the alias name of the OCS.
<i>host_name</i>		Specifies the host name or IPv4/IPv6 address of the OCS, such as www.bluecoat.com .

Table 16–3 HTTPS Origination Commands (Continued)

Option	Parameters	Description
https	[=port_number]	Specifies the port number on which the OCS is listening.
server		Specifies to use the relative path for URLs in the HTTP header because the next hop is a Web server, not a proxy server. Proxy is the default.
ssl-verify-server=	yes no	Specifies whether the upstream server certificate should be verified. You can only enable this command if the upstream host is a server, not a proxy.

The next scenario is useful when the ProxySG is deployed as a reverse proxy. This scenario is used when it is not necessary for a secure connection between the proxy and server. For information on using the ProxySG as a reverse proxy, see [Section D: "Selecting an HTTP Proxy Acceleration Profile"](#) on page 173.

Table 16–4 Scenario 2: HTTP Reverse Proxy with HTTPS Origination

HTTPS Reverse Proxy	HTTPS Origination
Client > HTTPS > ProxySG	ProxySG > HTTPS > Origin Content Server
Steps • Configure a keyring • Configure the SSL client • Configure the HTTPS service	Steps • Server URL rewrite -or- • Add a forwarding host (only for SGOS 3.1 or higher) • Set an HTTP port

Using server URL rewrite is the preferred method. For information on rewriting the server URL, refer to the *<Emphasis>Blue Coat Content Policy Language Guide*.

To configure HTTP origination:

At the (config) command prompt, enter the following commands:

```
SGOS#(config forwarding) create host_alias host_name
http[=port_number] server
```

where:

Table 16–5 HTTP Origination Commands

host_alias	alias_name	Specifies the alias name of the OCS.
host_name		Specifies the host name or IPv4/IPv6 address of the OCS, such as www.bluecoat.com .
http	[=port_number]	Specifies the port number on the OCS in which HTTP is listening.

Table 16–5 HTTP Origination Commands (Continued)

<code>server</code>		server specifies to use the relative path for URLs in the HTTP header because the next hop is a Web server, not a proxy server. Proxy is the default.

Creating Policy for HTTP and HTTPS Origination

Forwarding hosts must be already created on the ProxySG before forwarding policy can be created.

To create a policy using CPL:

```
<forward>
  url.host=host_name forward(host_alias)
```

To create a policy using VPM:

1. In the VPM, create a **Forwarding Layer**.
2. Set the **Destination** to be the URL of the OCS.
3. Set the **Action** to forward to the forwarding host and configure parameters to control forwarding behavior.

Chapter 17: Using the ProxySG in an IPv6 Environment

This section describes how a ProxySG can be configured to work in an IPv6 environment and explains the secure Web gateway and ADN features that support IPv6. It is assumed that you understand the underlying IPv6 technology.

Topics in this Section

This section includes information about the following topics:

- ❑ ["Using a ProxySG as an IPv6 Secure Web Gateway"](#)
- ❑ ["Using ProxySG Appliances in an IPv6 Application Delivery Network"](#)
- ❑ ["IPv6 Support on the ProxySG"](#) on page 337
- ❑ ["Configuring the ProxySG to Work in an IPv6 Environment"](#) on page 344
- ❑ ["Configuring an ADN for an IPv6 Environment"](#) on page 346
- ❑ ["Optimizing ISATAP Traffic"](#) on page 347
- ❑ ["Configuring IPv6 Global Settings"](#) on page 349
- ❑ ["IPv6 Policies"](#) on page 349

Using a ProxySG as an IPv6 Secure Web Gateway

The ProxySG's secure Web gateway functionality operates in both IPv4 or IPv6 networks. It provides visibility and control of all Web user communications — through authentication, authorization, logging, reporting, and policy enforcement — to create a productive, safe Web environment. The ProxySG has proxy support for multiple protocols and can control encrypted traffic for all users and applications inside and outside the enterprise. In addition, the ProxySG has built-in Web content filtering and content controls to prevent users from accessing inappropriate content using company resources.

In addition to its security and caching capabilities, the ProxySG offers functionality as an IPv4-to-IPv6 transition device. When an IPv6-enabled ProxySG is deployed between IPv4 and IPv6 networks, IPv4 clients will be able to access resources and services that are available only in the IPv6 domain. Likewise, IPv6 clients can access IPv4 resources when an IPv6-enabled ProxySG is part of the deployment. The ProxySG understands both IPv4 and IPv6 addresses, handles the DNS resolution of IPv4 and IPv6, and provides multiple proxy services that work in an IPv6 environment.

Using ProxySG Appliances in an IPv6 Application Delivery Network

Symantec's WAN optimization solution works in an IPv4, IPv6, or combination IPv4/IPv6 Application Delivery Network (ADN). If you are using Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) to transition a network from IPv4 to IPv6, the ProxySG can optimize this traffic as well. (See ["Using the ProxySG in an ISATAP Network"](#) on page 334.)

IPv6 is supported on the following types of ADN deployments:

- Open, unmanaged ADN
- Managed ADN
- Transparent deployments
- Explicit deployments
- Transparent load balancing
- Explicit load balancing

For information on configuring ADN for IPv6, see ["Configuring an ADN for an IPv6 Environment"](#) on page 346.

ADN-Managed Networks


In an ADN-managed network, the primary and backup ADN managers can be either IPv4 or IPv6, managed nodes can connect to the manager using IPv4 or IPv6, and the manager can advertise IPv4 and IPv6 routes. However, only IPv4 routes are advertised to managed nodes running older (pre-6.2.4) versions of SGOS.

To determine whether an ADN tunnel is IPv4 or IPv6:

1. Log in to the Management Console of the Concentrator peer.
2. Select **Statistics > Sessions > Active Sessions > ADN Inbound Connections**.
3. Locate the ADN peer address to see whether its format is IPv4 or IPv6. For example, if the peer address is `2001:418:9804:111::169`, it is an IPv6 tunnel. Or if the peer address is `10.9.45.129`, it is an IPv4 tunnel.

Transparent Deployment


In a transparent deployment, the Concentrator is installed physically in-path, and can intercept IPv4 and IPv6 connections from Branch peers.



Explicit Deployments

In explicit deployments, each Concentrator peer advertises the IPv4 or IPv6 server subnets that it fronts. The Concentrator peer can also act as an Internet gateway for IPv4 and IPv6 addresses, and subnets that are exempt from the Internet gateway can be IPv4 or IPv6. The explicit tunnel between the Branch and

Concentrator peers can be IPv4 or IPv6. A concentrator will be chosen as an Internet gateway only if it advertises at least one interface IP of the same address family as the destination server address.



Transparent Load Balancing

The ProxySG can intercept IPv4 or IPv6 connections and load balance either type of connection in a connection forwarding cluster. The cluster can contain both IPv4 and IPv6 addresses. The connection forward IP can be of the same type (for example, IPv6 for IPv6, IPv4 for IPv4) or a different type (IPv6 for IPv4 or vice versa) as the incoming connection. All ProxySG appliances in the forwarding cluster must be able to handle the address type of the connection.

Explicit Load Balancing


When using explicit tunnels, you can load balance for IPv4 and IPv6. When configuring load balancing with server subnets, multiple Concentrator peers can front an IPv4/IPv6 subnet. If multiple Concentrator peers are configured as Internet gateways, Branch peers will choose only those Concentrator peers that contain at least one address of the same family as the destination address.

When using an external load balancer, you can configure an IPv4 or IPv6 external virtual IP (VIP) address. The VIP address type (IPv4 vs. IPv6) must be reachable from all Branch peers.

Using the ProxySG in an ISATAP Network

One way to transition a network from IPv4 to IPv6 is with the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP). ISATAP uses a tunneling approach to transport IPv6 traffic across an existing IPv4 infrastructure by

encapsulating IPv6 packets with an IPv4 header. ISATAP-based connectivity can immediately be used to deliver IPv6 services while the IPv4-only infrastructure is gradually migrated to integrate native IPv6 capabilities. The tunneling of IPv6 traffic through the use of IPv4 encapsulation is called *6-in-4*.




In this example of an ISATAP topology, remote IPv6 clients need to access IPv6 servers over the enterprise IPv4 network. To accomplish this, IPv6 traffic from the client is encapsulated by the ISATAP router before traversing the IPv4 network. For example, IPv6 packets destined for IPv6 Server 1 in the data center are encapsulated with the IPv4 tunnel address of ISATAP Tunnel 1. IPv6 packets destined for the Internet are encapsulated with the IPv4 tunnel address of ISATAP Tunnel 2.

How Does the ProxySG Handle ISATAP Traffic?

After the ProxySG appliance identifies ISATAP traffic, it identifies the service inside the encapsulated packet, then uses the appropriate proxy to optimize the traffic. For example, the HTTP proxy optimizes web traffic with object caching, byte caching, compression, TCP optimization, and protocol optimization (assuming an ADN peer is found). For non-TCP, non-UDP, and services that are not intercepted (such as ICMPv6), the ProxySG uses the ISATAP proxy; this proxy

optimizes the IPv6 packet and payload using byte caching and compression over an ADN tunnel (assuming a peer is found). The following flow diagram describes how the ProxySG processes ISATAP traffic.



Notes:

- ❑ If the requested object is in cache or if the security policy determines that the request should not be allowed, the response is sent back to the client immediately over the encapsulated client-side connection.
- ❑ ISATAP is disabled by default.
- ❑ Reflect Client IP settings do not apply to the outer encapsulation header (the IPv4 address). Reflect Client IP settings are honored only for inner IPv6 source addresses for connections intercepted by application proxies, not the ISATAP proxy.

Feature Requirements

- ❑ The routers must support ISATAP.
- ❑ The ProxySG appliances must be inline between the ISATAP-capable routers.
- ❑ When load balancing is done via an external VIP, the concentrator should have SGOS 6.4 or higher.
- ❑ ISATAP must be enabled. See "[Optimizing ISATAP Traffic](#)" on page 347.

Feature Limitations

- ❑ Features that modify the destination address, such as URL rewrites and advanced forwarding, can cause issues with ISATAP processing because the IP encapsulation information must be preserved. If the destination address gets modified, users will see TCP connection errors because the server cannot be found.
- ❑ Only explicit ADN deployments are supported for ISATAP encapsulated traffic. The ProxySG uses the destination address in the encapsulation header to perform the route lookup for establishing the explicit ADN tunnel.
- ❑ In a virtually inline (WCCP) deployment, the When load balancing is done via an external VIP, the concentrator should have SGOS 6.4 or higher. is able to handle the ISATAP traffic and optimize the services for which application proxies are available, but the ISATAP proxy is not able to optimize the remaining ISATAP traffic, as it can in an inline deployment. This limitation occurs because the remaining traffic will likely not be redirected to the ProxySG.

IPv6 Support on the ProxySG

The ProxySG offers extensive support for IPv6, although there are some limitations. For details, see the following sections:

- ❑ ["IPv6 Proxies"](#) on page 337
- ❑ ["ISATAP Proxy"](#) on page 338
- ❑ ["ProxySG Management Access over an IPv6 Network"](#) on page 339
- ❑ ["Features that Support IPv6"](#) on page 340
- ❑ ["IPv6 Limitations"](#) on page 344
- ❑ ["Related CLI Syntax for IPv6 Configuration"](#) on page 344

IPv6 Proxies

The following proxies have underlying protocols that support IPv6 and can communicate using either IPv4 or IPv6:

Table 17–1


Proxy	For More Information
DNS	"Managing the Domain Name Service (DNS) Proxy"
FTP	"Managing the File Transport Protocol (FTP) Proxy"
HTTP	"Intercepting and Optimizing HTTP Traffic"
HTTPS	"Configuring and Managing an HTTPS Reverse Proxy"
SOCKS	"Managing a SOCKS Proxy"
SSL	"Managing the SSL Proxy"
RTSP	"Managing Streaming Media"

Table 17–1

Proxy	For More Information
TCP Tunnel	
Telnet Shell	"Managing Shell Proxies"

ISATAP Proxy

When the ProxySG encounters Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) traffic, it decides whether to process the 6-in-4 packets with the ISATAP proxy or one of the traditional application proxies (HTTP, FTP, CIFS, etc.). To make the decision on which proxy to use, the ProxySG identifies the service inside the encapsulated packet. If the ProxySG is intercepting this service, the traffic is processed by one of the traditional application proxies. If the service is not intercepted, the ProxySG uses the ISATAP proxy to optimize the IPv6 packet and payload over an ADN tunnel, assuming an ADN peer is found. Note that this proxy processes all ISATAP traffic that is not handled by application proxies, including ICMP, UDP, TCP, and routing protocols. If an ADN peer is not found, the packet cannot be optimized; it is simply sent to its destination.



The ISATAP proxy uses the following techniques to optimize the IPv6 packets:

- ❑ Byte caching
- ❑ Compression

The ISATAP proxy works differently than the application proxies: it processes individual packets instead of entire streams. It does not inspect the contents of the payload; it optimizes the entire packet.


Traffic that is processed by the ISATAP proxy appears in Active Sessions as the ISATAP_tunnel service and the ISATAP proxy type. The Active Sessions report lists the IPv4 tunnel address (not the IPv6 destination) as the server address since the ISATAP proxy has no insight into the payload of the packet.

The ISATAP proxy is not enabled by default. Until you enable ISATAP, 6-in-4 packets will be bypassed. See ["Optimizing ISATAP Traffic"](#) on page 347.

ProxySG Management Access over an IPv6 Network

All management services are available over IPv6 connectivity. If you have already defined IPv4 and IPv6 addresses for each ProxySG interface, both or either of these addresses can be selected as listeners for the HTTP-Console, HTTPS-Console, SSH-Console, and Telnet-Console services. The default setting for the service listener, **All SG IP addresses**, indicates that the management service is capable of accepting both IPv4 and IPv6 connections. When specifying IPv6 addresses, only global (not link-local) addresses can be used.

Use the **Configuration > Services > Management Services** option to view or modify the listeners for each management service.



To access the management console over a secure IPv6 connection, open a Windows Explorer or Firefox browser and enter the following in the address line:

```
https://[<ipv6 address>]:8082
```

where *<ipv6 address>* is the IP address that conforms to IPv6 syntax. Note that the square brackets must surround the IPv6 address when a port number is specified. For example:

```
https://[2001:db8::1]:8082
```

Note that Firefox has a bug that requires a backslash after the ending square bracket. For example:

```
https://[2001:db8::1]\:8082
```

You can also specify a host name that resolves to an IPv6 address. In this case, no brackets are required. For example:

```
https://atlantis:8082
```

To access the ProxySG using SSH, specify the IPv6 address enclosed in square brackets, for example:

```
[2001:db8::1]
```

For Telnet access, the square brackets are not required, for example:

```
2001:db8::1
```

Features that Support IPv6

The SGOS software accommodates the entry of either IPv4 or IPv6 IP addresses in applicable features. Table 17–2 lists the features that can be configured with either IPv4 or IPv6 addresses.

Table 17–2

Feature	For More Information	CLI Example
Network Interfaces	"Configuring a Network Adapter" on page 1237	<pre>#(config interface 0:0)ip-address 2001:db8::1428:57ab</pre>
DNS servers	"Adding DNS Servers to the Primary or Alternate Group" on page 816	<pre>#(config dns forwarding)edit primary #(config dns forwarding primary)add server 2001:db8:85a3::8a2e:370:7334</pre>
NTP servers	"Synchronizing to the Network Time Protocol" on page 40	<pre>#(config) ntp server 2001:db8::1428:57ab 8081</pre>
Default gateways	"Switching to a Secondary Default Gateway" on page 797	<pre>#(config)ip-default-gateway fe80::1%0:0 primary 100</pre>
Management service listeners	"Creating a Management Service" on page 1268	<pre>#(config management-services)edit HTTP- Console #(config HTTP-Console)add 2001:db8::1428:57ab 8081</pre>
Proxy service listeners	"Creating Custom Proxy Services" on page 120	<pre>#(config proxy-services)edit ftp #(config FTP)add all 2001::1/128 21 intercept</pre>
Static bypass entries	"Adding Static Bypass Entries" on page 143	<pre>#(config proxy-services) static-bypass #(config static-bypass) add 1000::/64 all</pre>
Restricted intercept lists	"About Restricted Intercept Lists" on page 147	<pre>#(config proxy-services) restricted- intercept #(config restricted-intercept) add all 2001::/64</pre>
Static routes	"Defining Static Routes" on page 800	<pre>#(config)inlinestatic-route-table eof 2000::/64 fe80::1%0:0 2001::/64 fe80::2%0:1 eof</pre>

Table 17–2

Feature	For More Information	CLI Example
Forwarding hosts	"Creating Forwarding Hosts" on page 875 "IPv6 Forwarding" on page 342	#(config forwarding) create host ipv6-proxy 2001:db8::1 http proxy
ADN managers	"Configuring the ADN Managers and Enabling ADN" on page 726	#(config adn manager) primary-manager 2001:418:9804:111::169 #(config adn manager)backup-manager 2001:418:9804:111::168
ADN server subnets	"Advertising Server Subnets" on page 729	#(config adn routing server-subnets) add 2001:418:9804:111::/64
ADN exempt subnets	"Configuring an ADN Node as an Internet Gateway" on page 750	# (config adn advertise-internet-gateway) exempt-subnets add 1234::/10
SMTP servers	"Enabling Event Notification" on page 1310	#(config smtp) server 2001:db8::1428:57ab
Syslog servers	"Syslog Event Monitoring" on page 1312	#(config event-log)syslog add 2001:418:9804:111::168
Load Balancer VIP	"Configuring Explicit Load Balancing Using an External Load Balancer" on page 748	#(config adn load-balancing) external-vip 2001:418:9804:111::200
Health checks on IPv6 hosts	"Creating User-Defined Host and Composite Health Checks" on page 1387	#(config health-check) create tcp ipv6-host-check 2001:db8::1 8080
Upload archive configurations on IPv6 servers	"Creating and Uploading an Archive to a Remote Server" on page 90	#(config) archive-configuration host 2001:db8::1
Upload access logs to an IPv6 server	"Editing Upload Clients" on page 635	#(config log log_name) ftp-client primary host 2001:418:9804:111::200
VPM objects	"VPM Objects that Support IPv6" on page 343	n/a
CPL objects	"CPL Objects that Support IPv6" on page 343	n/a

The IP address or hostname fields for these features accommodate the entry of IPv4 or IPv6 addresses and, when applicable, include a field for entering the prefix length (for IPv6 addresses) or subnet mask (for IPv4 addresses).

New Bypass List Entry

Client address

☐ All clients

☒ Client host or subnet

IP Address: 1026::93

Prefix/Subnet: 64

Server address

☒ All servers

☐ Server host or subnet

IP Address:

Prefix/Subnet:

OK Cancel

IPv6 Forwarding

To minimize WAN traffic, you can create forwarding hosts — the ProxySG configured as a proxy to which certain traffic is redirected for the purpose of leveraging object caching. (See ["About the Forwarding System"](#) on page 867.) It is possible to create IPv4-to-IPv6 forwarding, IPv6-to-IPv4 forwarding, and IPv6-to-IPv6 forwarding.

For example, to create a policy that forwards an IPv4 destination to an IPv6 forwarding host:

1. Create an IPv4 virtual IP (VIP) address for the ProxySG.
2. Create a forwarding host entry using an explicit IPv6 address or a hostname that resolves into an IPv6 address.
3. Launch the Visual Policy Manager (VPM)—**Configuration > Policy > Visual Policy Manager**.
4. Create or select a **Forwarding Layer**.
5. In a new rule, create a destination object: **Destination IP Address/Subnet** and enter the IPv4 VIP.
6. In the same rule, create an action object: **Select Forwarding** and select the configured IPv6 forwarding host.
7. Install the policy.

VPM Objects that Support IPv6

The VPM objects that support IPv6 addresses are listed below.

Source Objects

- **Client IP/Subnet**
- **Proxy IP Address/Port**
- **User Login Address**
- **RDNS Request IP/Subnet**

Destination Objects

- **DNS Response IP/subnet**
- **Destination IP/subnet**

Action Objects

- **Reflect IP (proxy IP)**

CPL Objects that Support IPv6

The following CPL objects support IPv6 addresses:

- `authenticate.credential.address()`
- `cache_url.address`
- `client.address`
- `dns.request.address`
- `dns.response.aaaa`
- `log_url.address`
- `proxy.address`
- `request.header.header_name.address`
- `request.header.referer.url.address`
- `request.x_header.header_name.address`
- `server.url.address`
- `url.address`
- `url.domain`
- `url.host`
- `user.login.address`

IPv6 Limitations

IPv6 support on the ProxySG has the limitations described below.

- ❑ The following proxies do not currently have IPv6 support:
 - MMS streaming
 - CIFS
 - MAPI
- ❑ The ProxySG does not intercept link-local addresses in transparent mode since this deployment isn't practical; transparent link-local addresses will be bypassed.
- ❑ IPv6 is not supported in a WCCP deployment.

Related CLI Syntax for IPv6 Configuration

The following are some CLI commands related to IPv6 configuration:

```
SGOS#(config) ipv6 auto-linklocal {enable | disable}
```

After link-local addresses are generated for the ProxySG interfaces, they will stay configured until they are manually removed using the `no ip-address` command or until the ProxySG is rebooted.

```
SGOS#(config interface interface_number) ipv6 auto-linklocal {enable | disable}
```

Enables or disables the automatic generation of link-local addresses for this interface. After a link-local address is generated for an interface, it remains configured until it is manually removed using the `no ip-address` command or until the ProxySG is rebooted.

```
> ping6 {IPv6_address | hostname}
```

Use this command to verify whether an IPv6 host is reachable across a network.

```
> show ndp
```


Shows TCP/IP Neighbor Discovery Protocol (NDP) table. NDP performs functions for IPv6 similar to ARP for IPv4.

Configuring the ProxySG to Work in an IPv6 Environment

Symantec's implementation of IPv6 support requires minimal IPv6-specific configuration. IPv6 support is enabled by default.

To configure a ProxySG to work in an IPv6 environment:


1. Assign IPv4 and IPv6 addresses to each interface on the ProxySG. You can add a link-local or global IPv6 address to any interface. Select the **Configuration > Network > Adapters** tab.
2. Select an interface and click **Edit**. The Configure a Native VLAN dialog displays.



3. Assign addresses:
 - a. Click **Add IP**. The Add IP Address dialog displays.
 - b. Enter the IPv6 address.
 - c. Click **OK** twice to close each dialog.
 - d. Click **Apply**.
4. Add a DNS server for IPv6. Select the **Configuration > Network > DNS > Groups** tab.

Groups		
Imputing		
DNS Groups		
Group Name	Servers	Domains
primary	10.2.2.100	*
alternate	2002::6	*

5. You can place both network servers types (IPv4 and IPv6) in the same DNS group, or separate them into different groups.
 - a. Click **Edit** or **New** and add a DNS server for IPv6.
 - b. Click **Apply**.
6. IPv6 requires its own gateway. Select the **Configuration > Network > Routing > Gateways** tab.



7. Define two default gateways: one for IPv4 and one for IPv6:
 - a. Click **New**. The Add List Item dialog displays.
 - b. Create a gateway to be used for IPv6.
 - c. Click **OK** to close the dialog.
 - d. Click **Apply**.
 - e. Repeat Steps a - d to create an IPv4 gateway (if you haven't done so already).
8. (Optional) Create policy for IPv6 servers. See "[IPv6 Policies](#)" on page 349.

Configuring an ADN for an IPv6 Environment

In addition to performing the steps in "[Configuring the ProxySG to Work in an IPv6 Environment](#)" on page 344, you need to configure the ADN for IPv6. See [Chapter 32: "Configuring an Application Delivery Network"](#) on page 711 and follow the steps for configuring the deployment applicable to your situation.

When transitioning your application delivery network to IPv6, be aware of the following considerations. While you are in the process of upgrading your ADN nodes to a version that supports IPv6 on ADN, you will have some nodes with the support and others without. During this transition period, the tunnel between the Branch and Concentrator peers may use only IPv4. For example, when a Concentrator is running SGOS 6.2.4 (a version with IPv6 ADN support) and the Branch peer is running SGOS 6.2.3 (without IPv6 ADN support), the tunnel will use only IPv4; only IPv4 routes are advertised to managed nodes running pre-6.2.4 software. When a pair of Concentrator and Branch peers are running SGOS 6.2.4 or higher, they can form a tunnel that supports IPv6.

When upgrading to an SGOS release that supports IPv6 on ADN, you should upgrade nodes in the following order:

1. **Upgrade ADN Managers.** In a managed ADN, the ADN managers must be upgraded *before* the other ADN peers. The managers must be assigned both IPv4 and IPv6 addresses.
2. **Upgrade Concentrator Peers.** The Concentrators should be configured with both IPv4 and IPv6 addresses until all Branch peers have been upgraded and configured to have IPv6 connectivity to the Concentrators. For external explicit load balancing, the Concentrator peers must be configured with an IPv4 external VIP until all Branch peers have been configured for IPv6.
3. **Upgrade load balancing ProxySG.** Any ProxySG that is doing load balancing should be upgraded before other ProxySG appliances.
4. **Upgrade Branch Peers.**
5. After all the managed nodes have been upgraded, configure the nodes with the manager's IPv6 address. The manager can then be reconfigured with an IPv6-only address, if the managed nodes have to connect to the manager using IPv6. The managed nodes can connect to the manager using either IPv4 or IPv6.

Note: In a transparent deployment, if a Branch peer running SGOS with IPv6 ADN support attempts to make a transparent connection, the request will not be intercepted by a Concentrator running a version without IPv6 ADN support. This is why you should upgrade Concentrators before Branch peers.

Optimizing ISATAP Traffic

The ProxySG can see inside a 6-in-4 encapsulated packet so that it can identify the service and use the appropriate proxy to optimize the traffic. For example, the Flash proxy optimizes Flash streaming traffic with object caching, TCP optimization, and protocol optimization (assuming an ADN peer is found). For services that are not intercepted (such as ICMPv6 and UDP), the ProxySG uses the ISATAP proxy; this proxy optimizes the IPv6 packet and payload using byte caching and compression over an ADN tunnel (assuming a peer is found).

1. Make sure your ProxySG appliances are inline between ISATAP-capable routers.
2. Enable both ISATAP options in the CLI.
 - a. Access the ProxySG CLI, with enable (write) access.
 - b. Type `conf t` to go into configuration mode.
 - c. At the `#(config)` prompt, type the following CLI commands:

```
isatap adn-tunnel enable
isatap allow-intercept enable
```
3. Use the Active Sessions report to verify ISATAP traffic is being processed and optimized by the appropriate proxy: ISATAP or the applicable application proxy.

For Intercepted Services:

- The Service name and Proxy type listed for the session correspond to the applicable application proxy (for example, HTTP or CIFS)
- An IPv6 address is listed for the Server.
- Colored icons should appear for the acceleration techniques that are applicable to the proxy: Compression , Byte Caching , Object Caching , Protocol Optimization 

For Non-TCP, Non-UDP, and Bypassed Services:

- The Service name listed for the session is ISATAP_tunnel, and the Proxy type is ISATAP.
- An IPv4 address is listed for the Server.
- Colored icons should appear for Compression  and Byte Caching 


Configuring IPv6 Global Settings

For details on IPv6 support on the ProxySG, see ["IPv6 Support on the ProxySG"](#) on page 337 and ["Configuring the ProxySG to Work in an IPv6 Environment"](#) on page 344.

IPv6 support is enabled by default, meaning that the ProxySG processes incoming IPv6 packets.

To change IPv6 global settings:

1. From the Management Console, select **Configuration > Network > Advanced > IPV6**.



2. Configure the **IPv6 Settings**:
 - a. To bypass all IPv6 traffic, select **Enable IPv6 force-bypass**. When this is selected, all IPv6 traffic is bridged or routed.
 - b. To have the ProxySG route bypassed traffic, select the **Enable IPv6 forwarding** option. When this option is disabled (as it is by default), the ProxySG discards bypassed traffic that is processed at layer-3.
3. Click **Apply**.

See Also

- ❑ ["IPv6 Support on the ProxySG"](#) on page 337
- ❑ ["Configuring the ProxySG to Work in an IPv6 Environment"](#) on page 344
- ❑ ["Configuring an ADN for an IPv6 Environment"](#) on page 346
- ❑ ["IPv6 Policies"](#) on page 349

IPv6 Policies

With the global policy for DNS lookups, the ProxySG first uses the configured IPv4 DNS servers for processing DNS requests. If this lookup fails, the ProxySG looks up the host on the configured IPv6 DNS servers. This processing of DNS requests happens automatically. To change the global setting for IP connection type preference, use the following policy:

```
server_url.dns_lookup(dns_lookup_value)
```

where

```
dns_lookup_value = ipv4-only|ipv6-only|prefer-ipv4|prefer-ipv6
```


If you have a known list of servers that are on IPv6 networks, you can avoid timeouts and unnecessary queries by creating policy to look up host names on IPv6 DNS servers only. For example:

```
<Proxy>
url.domain=etrade.com server_url.dns_lookup(ipv6-only)
url.domain=google.com server_url.dns_lookup(ipv6-only)
```

This policy overrides the global policy and look up the specified hosts (etrade.com and google.com) on the IPv6 DNS servers only.

To create DNS lookup policy in the Visual Policy Manager (VPM):

1. Launch the VPM and add or edit a **Web Access Layer**.
2. Create a new **Destination** object for an IPv6 host.
3. Create an **Action** object: **Set Server URL DNS Lookup**. The Add Server URL DNS Lookup Object dialog displays.



4. Select **Look up only IPv6 addresses**.
5. Repeat steps 2-4 for each IPv6 host.
6. Click **Install Policy**.

Chapter 18: Client Geolocation

To comply with local regulations, assist with traffic analysis, or reduce the risk of fraud and other security issues, you may need to know the origin of traffic in your network. SGOS supports client geolocation in reverse proxy mode, in which you can identify the source of traffic through the appliance based on IP address (and when applicable, the *effective client IP address*; refer to the *Visual Policy Manager Reference* for information on effective client IP address).

Before using geolocation, you must download a database that maps IP addresses to the countries with which they are associated. It also provides the supported names and ISO codes for countries, such as “United States [US]”. You can use country names or codes in policy to perform actions such as denying traffic from specific countries.

Topics in this Chapter:

The following sections describe how to configure and use geolocation:

- ❑ ["Prerequisites for Using Geolocation"](#) on page 352
- ❑ ["Enable Geolocation"](#) on page 353
- ❑ ["Download the Geolocation Database"](#) on page 354
- ❑ ["Determine Locations of IP Addresses for Incoming Connections"](#) on page 356
- ❑ ["Troubleshoot Geolocation"](#) on page 358
- ❑ ["Access Log Errors"](#) on page 359
- ❑ ["Remove Geolocation Settings"](#) on page 360

Section 1 Prerequisites for Using Geolocation

Before you can use geolocation, you must:

- ❑ Verify that you have a valid license for the feature. In the Management Console, select **Maintenance > Licensing > View** and look for license details in the **Intelligence Service Bundles** section.

If you do not have a valid license, the appliance is unable to download the database and the Management Console might display Health Monitoring errors. The access logs might also display error messages about the subscription. Review "[Troubleshoot Geolocation](#)" on page 358 for more information.

- ❑ Enable the geolocation service. See "[Enable Geolocation](#)" on page 353.
- ❑ Download the geolocation database. See "[Download the Geolocation Database](#)" on page 354.

Section 2 Enable Geolocation

Before you can use geolocation features, you must enable the geolocation service on the appliance.

Enable geolocation:

1. In the Management Console, select **Configuration > Geolocation > General**.
2. On the General tab, select the **Enable Geolocation functionality on the device** check box.
3. Click **Apply**.

The appliance starts to download the geolocation database. Allow the download to complete before attempting to use geolocation features.

Note: Refer to the *Command Line Interface Reference* for the related CLI command for enabling geolocation.

See Also

- ["Determine Locations of IP Addresses for Incoming Connections"](#) on page 356
- ["Troubleshoot Geolocation"](#) on page 358

Section 3 Download the Geolocation Database

When you enable the geolocation service, the appliance starts to download the database in the background. The Management Console displays the download in progress; wait for the download to complete before attempting to use the feature. The License and Download Status section on the **Download** tab displays statistics when download is complete.


If necessary, you can manually initiate the download database updates.

Manually download the geolocation database:

1. In the Management Console, select **Configuration > Geolocation > General**.
2. On the **Download** tab, in the Download Options section, click **Download Now**. When the download starts, the section displays a “Download is in progress” message.

If you receive a download error, check your network configuration and make sure that the appliance can connect to the Internet.

If the download is successful, the License and Download Status section displays statistics.



You can now write policy using country name or country code as defined in the geolocation database. You will also be able to see the supported country names and codes:

- ❑ In the Management Console (**Configuration > Geolocation > General > General** tab).
- ❑ In output for the `#show geolocation countries` and `#(config geolocation) view countries` CLI commands.
- ❑ When you add geolocation objects in the Visual Policy Manager (VPM).

Note: Refer to the *Command Line Interface Reference* for the related CLI command for downloading the geolocation database.

Cancel a Database Download in Progress

To stop any download of the Geolocation database that is currently in progress (including a download initiated from the CLI), click **Cancel Download** in the Download Options section on **Configuration > Geolocation > General > Download**. The console displays a “Canceling download” dialog. When the download is canceled, the dialog message changes to “Download Canceled”.

Section 4 Determine Locations of IP Addresses for Incoming Connections

After you download a geolocation database, you can identify the country that is associated with a given IP address. You can then use the country code or country name in policy.

Identify a country by IP address:

1. In the Management Console, select **Configuration > Geolocation > General**.
2. On the General tab, in the **IP** field, enter a valid IP address.
3. Click **Locate**.

The country associated with the IP address, as defined in the geolocation database, displays next to the IP field.

The screenshot shows the 'General' tab of the 'Geolocation' configuration page. At the top are three tabs: 'General' (selected), 'DNS Lookup', and 'Download'. Below the tabs is a checkbox labeled 'Enable Geolocation functionality on the device' which is checked. Underneath is a section titled 'Geolocation Lookup'. Inside this section, there is a label 'IP Address:' followed by a text input field containing '128.66.0.0'. To the right of the input field, the text 'Italy [IT]' is displayed. Below the input field is a 'Locate' button. At the bottom of the section is a link that says 'Show list of countries in Geolocation database'.

To write policy about a specific country, use the country code specified in square brackets beside the country name. For example, the country code for Italy is "IT".

Note: IP address mappings to locations change over time, and periodic geolocation database updates reflect these changes. A changed IP address mapping can cause the following behavior:

- geolocation lookup results are different from a previous lookup
- geolocation policy is no longer working as expected

Use current lookup results to update policy as appropriate.

If the IP address you enter is not valid, an error appears. See "[Troubleshoot Geolocation](#)" on page 358 for more information.

4. (Optional) To view the list of countries in the geolocation database, click **Show list of countries in Geolocation database**. The list opens in a separate browser window.

Use Cases: Write Geolocation Policy

Refer to the following examples of writing geolocation policy for incoming connections.

The `client.address.country=<"country_name">` condition returns the country from which traffic originates, based on the client IP address. For detailed usage information on policy gestures, refer to the *Content Policy Language Reference*.

Use Case 1

You require policy to allow client connections only from North America. You can use the following CPL:

```
; only accept client connections from North America
<proxy>
  allow client.address.country=(US, CA)
  deny("Restricted location: $( x-cs-client-ip-country) ")
```

Use Case 2

You require policy to allow client connections only from North America, including proxied traffic as specified. You can use the following CPL:

```
; only accept traffic from North America with support for proxied traffic
; with client address in X-Forwarded-For
<proxy>
  client.effective_address("${request.header.X-Forwarded-For}")

<proxy>
  client.effective_address.country=(US, CA) ok
  deny("Restricted location: $( x-cs-client-effective-ip-country) ")
```

Note: In the VPM, the **Client Geolocation** object is available in the Source column in policy layers. Refer to the *Visual Policy Manager Reference* for information.

Section 5 Troubleshoot Geolocation

You might encounter the following errors. Refer to the specified troubleshooting steps.

Note: *This is not a valid IP*

Cause: You have entered an invalid IP address in the **IP** field when performing geolocation lookups.

Fix: Correct the IP address and look it up again.

Note: *The geolocation database is currently unavailable*

Cause: No geolocation database is installed.

Fix: Download the database. See "[Download the Geolocation Database](#)" on page 354.

Note: *This device does not have a valid geolocation license*

Cause: The appliance does not have a valid geolocation license.

Fix: Ensure that the following are true:

- Each appliance in your deployment has its own license.
- The licensing status for your appliance is in good health.

Warning: *The geolocation database is not installed*

This error appears in the browser window that opens when you click **Show list of countries in Geolocation database** under Geolocation Lookup. It can also appear in the CLI.

Cause: No geolocation database is installed.

Fix: Download the database. See "[Download the Geolocation Database](#)" on page 354.

Section 6 Access Log Errors

The following access log errors may appear in the access log if there is a problem with your subscription.

- ❑ The Geolocation subscription file is out of date
- ❑ Failed trying to get the subscription settings from the Geolocation subscription file
- ❑ Failed trying to download the Geolocation subscription file
- ❑ Failed trying to extract and activate the Geolocation payload file

Note: If you receive other errors while setting up or using geolocation, refer to MySymantec.

Section 7 Remove Geolocation Settings

To remove geolocation settings, purge the geolocation database and remove policy.

Remove the Database

Remove the database:

1. Disable geolocation.
 - a. In the Management Console, select **Configuration > Geolocation > General**.
 - b. On the General tab, select the **Enable Geolocation functionality on the device** check box.
 - c. Click **Apply**.
2. Log in to the CLI and enter the following command:

```
#(config geolocation)purge
```

The CLI returns to the `#(config geolocation) node`.

You can issue the **view countries** command to verify that the geolocation database has been removed. The output should show no list of countries and warn that the database is not installed:

```
#(config geolocation)purge
```

```
#(config geolocation)view countries
```

```
Countries defined by system:
```

```
Invalid
None
Unavailable
Unlicensed
```

```
Additional locations:
```

```
Countries defined by geolocation database:
```

```
Warning: The geolocation database is not installed
```

Remove Policy

Remove geolocation policy:

1. In the Management Console, select **Configuration > Policy > Visual Policy Manager**.
2. Click **Launch** to launch the VPM.
3. Remove any geolocation objects or rules. Refer to the *Visual Policy Manager Reference* for instructions.
4. Click **Install Policy**.
- 5.

Chapter 19: Web Application Protection

You can protect web applications from web attacks using Application Protection policy. Application security on the ProxySG appliance consists of the following:

- ❑ Application Protection, which is a component of Web Application Protections, a subscription-based offering that supplements the services available in Web Application Reverse Proxy. This feature downloads a SQL injection fingerprints database, which defines how the appliance detects attacks.
- ❑ Related content policy language (CPL), which you can write to prevent web attacks, including SQL and other types. Refer to the *Content Policy Language Reference* for more information.

Topics in this Chapter

The following sections describes Application Protection and how to use policy to protect your web applications from attacks.

- ❑ ["Using Application Protection" on page 362](#)
- ❑ ["Understanding the Risk Score" on page 365](#)
- ❑ ["Creating an Application Protection Solution" on page 367](#)
- ❑ ["Advanced Features for Web Application Protection" on page 373](#)
- ❑ ["Reference Information" on page 375](#)

Section A: Using Application Protection

Application Protection allows the ProxySG appliance to detect SQL injection attacks without the need to write and maintain SQL injection policy. When you enable the feature, it downloads a database containing the latest SQL injection fingerprints, culled from real-world attacks. This allows the appliance to detect numerous SQL injection attacks against different SQL databases that web applications use.

To keep the database up-to-date, you can configure settings to be notified whenever an update is available, or you can allow the service to automatically download new versions.

Prerequisite for Using Application Protection

Before you can set up Application Protection, you must have a valid license for it. Refer to your Sales Engineer for more information.

If you enable Application Protection but do not have a valid license, the Management Console and event log display errors indicating that the subscription could not be downloaded.

Enabling Application Protection

Before you can use Application Protection or related policy, you must enable the feature on the appliance. For information on using this feature in a test environment, see "[Testing the Application Protections](#)" on page 363.

To enable Application Protection:

1. In the Management Console, select **Configuration > Threat Protection > Application Protection**.
2. On the Application Protection tab, select the **Enable** check box.
3. Click **Apply**. The appliance attempts to download the database for the first time.

The service will automatically check for and download updates if:

- the service is enabled
- an Internet connection exists
- the notification setting (described in "[Testing the Application Protections](#)" on page 363) is disabled

What if the Initial Download is Not Successful?

If you receive a download error and the Management Console banner displays **Critical** shortly after you click **Apply**, the download might have failed. To confirm if this is the case, select **Statistics > Health Monitoring > Status** and look for the status "Application Protection failed on initial download" for **Subscription Communication Status**.

Note: The **Critical** error appears if the initial download attempt fails. After the database downloads successfully, the service periodically checks for a newer version of the database. If several update checks fail to connect to Symantec, a **Warning** error appears in Health Monitoring until the failure is corrected.

See Also

- ❑ ["Using Application Protection"](#) on page 362
- ❑ ["Verifying the Database Download"](#) on page 364
- ❑ ["About the Licensing Metrics"](#) on page 1341
- ❑ ["About the Status Metrics"](#) on page 1343

Testing the Application Protections

If you want to test the application protections before deploying them in a production environment, you can manually download fingerprint database updates to an appliance in your lab or staging area.

Enable Notification

To enable notification:

1. In the Management Console, select **Configuration > Threat Protection > Application Protection**.
2. (If feature is not already enabled) On the Application Protection tab, select **Enable**.
3. Select the **Do not auto download, but notify when a new version becomes available** check box.
4. Click **Apply**.

When a new database is available, a notification is sent to the administrator e-mail account and also recorded in the event log.

Download Updates Manually

To download the database manually:


1. In the Management Console, select **Configuration > Threat Protection > Application Protection**.
2. On the Application Protection tab, click **Download Now**.

See Also

- ❑ ["Verifying the Database Download"](#) on page 364

Verifying the Database Download

The License and Download Status field shows statistics about the previous successful and unsuccessful downloads. If the last download was unsuccessful, the field contains an error.



The screenshot shows the 'Application Protection' configuration page. It has a tabbed interface with 'Application Protection' selected. There are two checkboxes: 'Enable' (unchecked) and 'Do not auto download, but notify when a new version becomes available' (unchecked). Below these is the 'License and Download Status' section, which contains a text box with the following information:

```
Notify Only:           Disabled
Download method:       Direct
Unsuccessful downloads: 32
Last attempt:
  Time:                Fri, 30 Aug 2013 19:49:38 UTC
  Downloading from:     https://subscription.es.bluecoat.com/application-protection/database
  $ Failed to DNS resolve subscription.es.bluecoat.com
Last successful download:
  N/A
```

At the bottom of the section is a 'Download Now' button.

If you receive a download error, check your network configuration and make sure that the appliance can connect to the Internet.

Section B: Understanding the Risk Score

Web application security on the ProxySG appliance involves inspecting HTTP transactions for common web attacks and then weighing the risk of any attacks it detects. Policy does not trigger an action (such as denying the transaction) based on its risk assessment; rather, the appliance assigns the transaction a *risk score*, which represents the anomaly level of the transaction. You can write risk score policy to determine what action should be taken, if any, on the transaction.

Note: For risk score policy to take effect, you must enable Application Protection in the Management Console. See ["Using Application Protection"](#) on page 362. After you enable this feature, you can optionally write more web application protection rules. See ["Configuring Injection Protection"](#) on page 369 for instructions.

Although the appliance has a predefined default risk score for anomalous transactions, you can define risk tolerance on different web applications according to your needs using risk score policy.

Creating Actions Based on the Client's Risk Score

You can adjust the risk score-based triggers to set an action based on the cumulative risk score that is reached by a client. By default, each violation of SQL injection, invalid multipart form, null byte, HTTP parameter pollution, and multiple encoding adds a weight of 10 to the risk score.

Note: You cannot change the risk score weight for SQL injection violations, but you can change it for the other types of violations using content policy language (CPL).

Trigger an action based on the cumulative risk score.

```
risk_score=risk_score_value
```

where *risk_score_value* is an integer greater than or equal to zero. The default value is 10.

The threshold can be:

- An exact value, such as 10
- A range, such as 20..40 (a value between 20 and 40) or 10.. (a value equal to or greater than 10)

Example

The ProxySG appliance blocks the connection when it detects two or more web attacks in the same request; assume a risk-score value of 10 (the default).

```
<proxy>
  deny risk_score=20..
```

How Policy Inspects a Transaction and Calculates Cumulative Risk

The following is an example of how the appliance could inspect a transaction and calculate the risk score.

1. An attacker enters the query:

```
http://yourdomain.com/auth?Id=12345%20or%201=1
```

2. The ProxySG appliance detects the web attack.
3. The appliance normalizes the request. It analyzes the string name-value pairs in the request and (if needed) decodes URL-encoded and HTML-encoded data. In this example, `Id=12345%20or%201=1` is decoded as `Id=12345 or 1=1`.
4. Policy evaluates the relevant CPL:

```
; Scan for SQL injection attacks against all parts of  
; request (query string, cookie, and body)  
<proxy>  
    http.request.detection.injection.sql(yes)
```

```
; Block transactions where SQL injection is detected  
<proxy>  
    deny risk_score=10..
```

The appliance processes the query and detects an SQL injection attack pattern. The attack is counted as one violation with a score of 10. The policy above blocks the attack.

Section C: Creating an Application Protection Solution

To protect your web application from attacks that inject malicious content or otherwise change the behavior of your application, create an application protection solution that detects specific attack types.

This section describes the different web attacks that the ProxySG appliance can detect and some content policy language (CPL) examples that you can write for protection from such attacks.

- ❑ ["Preventing an SQL Injection Attack"](#) on page 368
- ❑ ["Configuring Injection Protection"](#) on page 369
 - ["Preventing a Null-Byte Injection Attack"](#) on page 369
 - ["Preventing an Invalid Multipart Form Attack"](#) on page 370
 - ["Preventing an HTTP Parameter Pollution Attack"](#) on page 370
 - ["Preventing a Multiple Encoding Attack"](#) on page 371
- ❑ ["Creating Custom Detection Rules"](#) on page 372
- ❑ ["Reference Information"](#) on page 375

About Injection Attacks

Injection attacks occur when malicious content is inserted into a field for execution. This type of attack takes advantage of a security vulnerability in the web application. Depending on the permissions configuration of a website, a successful attack can result in a completely or partially compromised database. The attacker can also gain access to confidential information such as personal records, as well as the ability to tamper with the database content.

The ProxySG appliance can detect certain injection attacks and conduct input validation in order to inspect requests, and, if appropriate, block suspicious transactions from being processed. The appliance can help prevent the following web attacks: SQL injection, null-byte injection, invalid multipart form content, HTTP parameter pollution (HPP), and multiple encoding.

To enable SQL injection protection, you must subscribe to the application protection subscription provided by the Symantec Web Application Protection service, which contains the latest SQL injection fingerprints.

Important Note about Web Application Protection and Request Body Size

By design, the ProxySG appliance only scans the first 8 KB of the request body. As a result, policy cannot detect malicious code beyond those first 8 KB. You can use the `http.request.body.size=` condition to test for requests larger than 8 KB and then prevent a bypass due to the position or volume of malicious content in a request.

Refer to the *Content Policy Language Reference* for information on the `http.request.body.size=` condition.

Preventing an SQL Injection Attack

To detect SQL injection attacks, the ProxySG appliance relies on the Application Protection database described in ["Using Application Protection"](#) on page 362.

Configuring SQL Injection Protection

To enable SQL injection protection on the ProxySG appliance, use the `http.request.detection.injection.sql` action in CPL. The object is available in the `<Proxy>` layer and is applied on a per-transaction basis.

Enable SQL injection protection and examine traffic for potentially malicious code.

```
<proxy>
http.request.detection.injection.sql[<attribute>, ...] (yes|no)
```

where:

<code><attribute></code>	Specifies the content to inspect: • <code>query</code> - Name and value in query string; unnamed values. • <code>cookie</code> - Name and value in Cookie and Cookie2. • <code>body</code> - Inspect up to the first 8 KB of request body data for SQL injection pattern. If none is specified, all sources are inspected.
<code>yes no</code>	Sets the action for the defined attribute(s).

Examples

Enable SQL injection protection to scan all inputs but do not inspect cookies from a specific domain.

```
<proxy>
http.request.detection.injection.sql (yes)
```

```
<proxy>
url.host=specialdomain.com \
http.request.detection.injection.sql[cookie] (no)
```

Enable SQL injection protection and inspect cookies and query string.

```
<proxy>
http.request.detection.injection.sql[cookie,query] (yes)
```

To prevent an SQL injection attack, write the following CPL:

```
; Scan for SQL injection attacks
<proxy>
http.request.detection.injection.sql (yes)

; Block transactions where SQL injection is detected
; if the cumulative risk score is 10 or higher
<proxy>
deny risk_score=10..
```

Configuring Injection Protection

To enable null-byte, HTTP parameter pollution (HPP), invalid multipart form, and multiple encoding attack detection on the ProxySG appliance, use the `http.request.detection.other` action in CPL. The object is available in the `<Proxy>` layer and is applied on a per-transaction basis.

Enable injection protection and examine traffic for potentially malicious code.

```
<proxy>
  http.request.detection.other[.<attribute>,...] (yes|no)
```

where:

<code><attribute></code>	Specifies the attack type: <code>null_byte</code> - Detects and sanitizes content that contains a null byte. <code>invalid_form_data</code> - Scans the request for invalid form content. <code>parameter_pollution</code> - Scans the request to find multiple instances of a parameter with the same name. <code>multiple_encoding</code> - Scans the request for multiple encoding attacks.
<code>yes no</code>	Sets the action for the defined attribute.

See Also

- ❑ ["Preventing a Null-Byte Injection Attack" on page 369](#)
- ❑ ["Preventing an Invalid Multipart Form Attack" on page 370](#)
- ❑ ["Preventing an HTTP Parameter Pollution Attack" on page 370](#)
- ❑ ["Preventing a Multiple Encoding Attack" on page 371](#)

Preventing a Null-Byte Injection Attack

A null-byte attack exploits the fact that null-byte string terminators are handled inconsistently amongst different programming languages. In this type of attack, the attacker embeds a null byte at a specific position in a string, causing the rest of the string to terminate.

To detect null-byte injection attacks, the ProxySG appliance scans query strings for URL-encoded null bytes, and more.

Null-Byte Injection Attack Prevention Policy

To prevent a null-byte injection attack, write the following CPL.

```
; Scan for null-byte injection attacks
<proxy>
    http.request.detection.other.null_byte(yes)

; Block transactions where null-byte injections are detected
; if the cumulative risk score is 10 or higher
<proxy>
    deny risk_score=10..
```

Preventing an Invalid Multipart Form Attack

The HTML form content type `multipart/form-data` is used for submitting forms that contain binary data or web form data of arbitrary size. A `multipart/form-data` form submission consists of a series of parts; each part represents a control type (such as button, menu, text, and file) that is used for selecting or entering data. When a user selects or enters all of the data, the parts are sent for processing to the user agent in the specified order.

To detect malicious data sent via multipart form submissions, the ProxySG appliance validates form boundary, form header format, header fields and form data. It also detects multiple instances of name and filename fields.

Invalid Multipart Form Attack Prevention Policy

To prevent an invalid multipart form attack, write the following CPL.

Example

```
; Scan for invalid multipart form data
<proxy>
    http.request.detection.other.invalid_form_data(yes)

; Block transactions where invalid multipart form is detected
; if the cumulative risk score is 10 or higher
<proxy>
    deny risk_score=10..
```

Preventing an HTTP Parameter Pollution Attack

An HTTP parameter pollution (HPP) attack occurs when an attacker injects encoded query string delimiters to HTTP GET or POST parameters to facilitate the addition of more HTTP parameters. These additional HTTP parameters can override existing ones to bypass input validation and inject malicious payloads.

The impact of HPP attacks depends on how the backend parses query strings. For example, if an attacker injects multiple values for the same name, PHP would use the last instance, JSP would use the first one, and ASP would concatenate the values. An attacker could analyze a web application to determine how the backend parses the query, whether parameters are not sanitized sufficiently, and then exploit the vulnerability.

To detect possible HPP attacks, the ProxySG appliance checks for names with multiple values, and more.

HTTP Parameter Pollution Attack Prevention Policy

To prevent an HPP attack, write the following CPL.

```
; Scan for polluted HTTP parameters
<proxy>
    http.request.detection.other.parameter_pollution(yes)

; Block transactions where HTTP parameter pollution is detected
; if the cumulative risk score is 10 or higher
<proxy>
    deny risk_score=10..
```

Preventing a Multiple Encoding Attack

Multiple encoding attacks exploit the fact that characters can be expressed in different formats; for example, in URL encoding, a backslash (“\”) is %5c.

Because a back-end server decodes user input only once, an attacker could enter characters that have been encoded more than once, for instance in path names or query strings, in order to evade security filters.

For example, URL encoding a backslash once results in:

```
%5c
```

Encoding %5c results in:

```
%255c
```

If an attacker enters %255c in a path name, the backend decodes it once, as %5c.

To detect multiple encoding attacks, the ProxySG appliance scans for characters in request parameters that appear to be still encoded after normalization.

Multiple Encoding Attack Prevention Policy

To prevent a multiple encoding attack, write the following CPL.

```
; Scan for URL-encoded characters in requests
<proxy>
    http.request.detection.other.multiple_encoding(yes)

; Block transactions where multiple encoding attack is detected
; if the cumulative risk score is 10 or higher
<proxy>
    deny risk_score=10..
```

Creating Custom Detection Rules

You can create custom rules based on your specific needs. Creating a custom rule allows policy to match against every name/value of every parameter in a request. The object is available in the <Proxy> layer and is applied on a per-transaction basis.

Create and enable custom detection rules that examine traffic for potentially malicious content.

```
<proxy>
  http.request[<attribute>,...].<modifier>[.case_sensitive]=pattern
```

where:

<attribute>	A comma-separated list of the predefined content sources. Refer to the <i>Content Policy Language Reference</i> for detailed information.
<modifier>	Specifies an expression type. The following options are available: • exact • substring • suffix • regex • count
case_sensitive	Specifies the test as case-sensitive. By default, the test is case-insensitive. Does not apply to the count modifier.

Examples

The following is an example using a regular expression.

```
; Perform regex scan for case-insensitive pattern "bad"
; in any name or value in cookie headers and cookie names
<proxy>
  http.request[cookie_name,cookie].regex="bad"
```

The following is an example using a case-sensitive regular expression.

```
; Perform regex scan for case-sensitive string "bad"
; in any name or value in any parameter
<proxy>
  http.request[name,value].regex.case_sensitive="bad"
```

Section D: Advanced Features for Web Application Protection

The default web application protection policy on the ProxySG appliance should be appropriate in most deployments; however, you can create policy that goes beyond the default settings to address your specific security and business requirements.

Specifying a Risk Score for a Web Attack Type

For each type of web attack, you can adjust the risk score-based trigger to set an action based on the cumulative risk score that a client reaches. Use the CPL:

```
risk_score.other[.<attribute>,...](risk_score_value)
```

where:

- **<attribute>** is the attack type (`null_byte`, `invalid_form_data`, `parameter_pollution`, or `multiple_encoding`)
- ***risk_score_value*** is the custom risk score you want to specify

Examples

```
; Set the risk score for null-byte attacks to 20
<proxy>
    risk_score.other.null_byte(20)

; Set the risk score for HPP attacks to 30
<proxy>
    risk_score.other.parameter_pollution(30)
```

Setting a Maximum Risk Score

To optimize how the ProxySG appliance processes malicious user data, you can specify that the appliance block requests and discontinue scanning requests immediately after the maximum allowable risk score is reached. Specifying a maximum helps decrease the load on appliance resources dedicated to processing web attacks.

Set a maximum risk score:

```
risk_score.maximum(integer)
```

where `integer` is the maximum risk score value before blocking requests and discontinuing scan requests.

You can specify an integer between 0 and 2147483647. The default value is 40. Specify 0 to never block requests due to a maximum risk score.

Example

The following example enables SQL injection attack detection and defines a maximum risk score. The appliance scans a request until it detects two anomalies, after which it blocks the request and evaluates the next policy rule. Assume a default risk score of 10.

```
; Scan for SQL injection attacks
<proxy>
    http.request.detection.injection.sql(yes)

; Block and discontinue scan requests when two anomalies are detected
<proxy>
    risk_score_maximum(20)

; Block transactions where SQL injection is detected
; if the cumulative risk score is 10 or higher
<proxy>
    deny risk_score=10..
```

Section E: Reference Information

This section contains reference information pertaining to web application protection.

Access Log Fields

Access log entries are generated upon detection of suspicious data when the following ELFF log formats are used:

`x-risk-score`

`x-risk-category`

Enable Access Logging from **Configuration > Access Logging > General** in the Management Console.

Substitutions

The total risk score and the category can be used as substitutions. Using this substitution allows you to fetch the values from the transaction and substitute them in a string. The following substitutions are available:

`$(risk-score)`

`$(risk-category)`

Chapter 20: Filtering Web Content

Content Filtering allows you to categorize and analyze Web content. With policy controls, content filtering can support your organization's Web access rules by managing or restricting access to Web content and blocking downloads from suspicious and unrated Web sites, thereby helping protect your network from undesirable or malicious Web content.

The ProxySG supports Symantec WebFilter as well as other third-party databases. This chapter describes how to configure the ProxySG to process client Web requests and to control and filter the type of content retrieved.

For information on integrating your local ProxySG Appliance content filtering policy with Symantec Cloud Service policy, please see *Universal Policy: Applying Global Policy to Local and Remote Users*.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ Section A: "Web Content Filtering Concepts"
- ❑ Section B: "Setting up a Web Content Filter"
- ❑ Section C: "Configuring Symantec WebFilter and WebPulse"
- ❑ Section D: "Configuring a Local Database"
- ❑ Section E: "Configuring Internet Watch Foundation"
- ❑ Section F: "Configuring a Third-Party Vendor"
- ❑ Section G: "About Blue Coat Categories for YouTube"
- ❑ Section H: "Applying Policy"
- ❑ Section I: "Troubleshooting"

Section A: Web Content Filtering Concepts

Content filtering is a method for screening access to Web content. It allows you to control access to Web sites based on their perceived content. On the ProxySG, using a content filtering database in conjunction with policy allows you to manage employee access to Web content and to restrict access to unsuitable content. Restricting access or blocking Web content helps reduce the risk of malware infections caused by visiting questionable sites.

This section discusses content filtering databases and categories, and the content filtering options available on the ProxySG.

About Content Filtering Categories and Databases

Content filtering categories comprehensively classify the vast and constantly growing number of URLs that are found on the Web into a relatively small number of groups or categories. These categories then allow you to control access to Web content through policy.

A content filtering database has a pre-defined set of categories provided by the content filtering vendor. Individual content filter providers such as Symantec WebFilter, define the content-filtering categories and their meanings. Depending on the vendor, a URL is listed under one or more categories. Each URL can support a maximum of 16 categories.

A content filtering database does not block any Web site or category. The role of the database is to offer additional information to the proxy server and to the administrator about a client request. After you configure your content filter provider and download the database, you can map the URLs to the list of categories. You can then reference these categories in policy and limit, allow, or block requests. Client access to a Web request depends on the rules and policies that you implement in accordance with company standards.

Some policies that you might create, for example, are as follows:

- ❑ Block or unblock specific sites, categories, or specific file types such as executables.
- ❑ Apply different filtering policy for each site or group within your organization, by IP address or subnet. If you wish to use password authentication to grant or deny access to the requested content, you must have configured authentication realms and groups on the ProxySG. For information on configuring authentication, see "[Controlling User Access with Identity-based Access Controls](#)" on page 898.
- ❑ Allow schedule-based filtering to groups within your organization.

A valid vendor subscription or license is required to download a content filter database. For example, Symantec WebFilter is licensed while some supported third-party vendors require a subscription.

If your subscription with the database vendor expires or if the available database is not current, the category **unlicensed** is assigned to all URLs and no lookups occur on the database. To ensure that the latest database version is available to you, by default, the ProxySG checks for database updates once in every five minutes.

About Application Filtering

In addition to URL category filtering, you can filter content by Web application and/or specific operations or actions done within those applications. For example, you can create policy to:

- ❑ Allow users to access all social networking sites, except for Facebook. Conversely, block access to all social networking sites except for LinkedIn.
- ❑ Allow users to post comments and chat in Facebook, but block uploading of pictures and videos.
- ❑ Prevent the uploading of videos to YouTube, but allow all other YouTube operations such as viewing videos others have posted. Conversely, preventing uploading but block access to some videos according to the video's category.
- ❑ Allow users to access their personal email accounts on Hotmail, AOL Mail, and Yahoo Mail, but prevent them from sending email attachments.

This feature allows administrators to block actions in accordance with company policy to avoid data loss accidents, prevent security threats, or increase employee productivity.

See ["Creating Policy for Controlling Web 2.0 Applications"](#) on page 430.

About the Content Filtering Exception Page

Exception pages are customized Web pages (or messages) sent to users under specific conditions defined by a company and its security policies. An exception page is served, for example, when a category is blocked by company policy.

The ProxySG appliance offers multiple built-in exception pages that can be modified to meet your enterprise needs. For content filtering, the ProxySG includes the `content_filter_denied` and `content_filter_unavailable` built-in exception pages.

The `content_filter_denied` exception page includes the following information:

- an exception page message that includes the content filtering category affecting the exception.
- (Only for Symantec WebFilter) A category review URL, where content categorizations can be reviewed and/or disputed.
To add the link in the message, select the checkbox **Enable category review message in exceptions** in **Configuration > Content Filtering > General**. See ["Enabling a Content Filter Provider"](#).

The `content_filter_unavailable` exception page includes a message that states the reason for the denial and provides a probable cause—the request was denied because an external content filtering service was not available owing to transient network problems, or a configuration error.

See "Applying Policy" for information on using the exception pages in policy.

For customizing the exception page, refer to the Advanced Policy Tasks chapter, Section E, of the *Visual Policy Manager Reference*.

Web Content Filtering Process Flow

The following diagram illustrates the process flow when Web content filtering is employed in the network. This diagram does not include the dynamic categorization process, for details on dynamic categorization, see "About the Dynamic Categorization Process" on page 385.




Figure 20–1 Web Content Filtering Process Flow

Supported Content Filter Providers

The ProxySG appliance supports several content filter providers. From the following options, you can use up to four URL content filters in any combination:

- ❑ Blue Coat WebFilter. Blue Coat WebFilter provides both an on-box content filtering database and the WebPulse service, a cloud-based threat-protection feature.
- ❑ Local database. Create and upload your custom content filtering database to the ProxySG appliance. This database must be in a text file format.
- ❑ The Internet Watch Foundation (IWF) database. For information about the IWF, visit their Web site at: <http://www.iwf.org.uk/>
- ❑ A supported third-party content filtering vendor database— Proventia or Optenet. You cannot use two third-party content filtering vendors at the same time.

Note: You cannot configure the legacy content filtering vendors, such as Intersafe, I-Filter, and Surfcontrol, using the Management Console. You must use the Command Line Interface (CLI) to modify configuration settings for these vendors. Refer to the *Command Line Interface Reference* for the list of CLI commands available.

If you are upgrading from an earlier SGOS version, and one of the vendors listed above is enabled on your ProxySG appliance, the message “You must use CLI to configure <vendor> content filtering” displays in the Management Console.

- ❑ YouTube. The appliance obtains video categories from the YouTube Data API v3.0. After you enable Blue Coat categories for YouTube, you can reference these categories in policy to control YouTube traffic.

Note: In April 2015, Google deprecated YouTube Data API v2.0. As a result, Blue Coat categories for YouTube are no longer supported in SGOS versions prior to version 6.5.7.6.

Note: This feature is provided on an "as-is" basis. Symantec has no control of, and is not responsible for, information and content provided (or not) by YouTube. Customer is required to apply and use its own API key in order to activate this feature, and therefore obligated to comply with all terms of use regarding the foregoing (for example, see <https://developers.google.com/youtube/terms>), including quotas, restrictions and limits on use that may be imposed by YouTube. Symantec shall not be liable for any change, discontinuance, availability or functionality of the features described herein.

About Symantec WebFilter and the WebPulse Service

Symantec WebFilter, in conjunction with the WebPulse service, offers a comprehensive URL-filtering solution. Symantec WebFilter provides an on-box content filtering database and WebPulse provides an off-box *dynamic categorization service* for real-time categorization of URLs that are not categorized

in the on-box database. WebPulse dynamic categorization includes both traditional content evaluation, for categories such as pornography, as well as real-time malware and phishing threat detection capabilities. WebPulse services are offered to all customers using Symantec WebFilter.

WebPulse is a cloud service that allows inputs from multiple enterprise gateways and clients and creates a computing grid. This grid consists of Symantec WebFilter, K9, and ProxyClient customers, who provide a large sample of Web content requests for popular and unrated sites. Based on the analysis of this large volume of requests, the computing grid continuously updates the master Blue Coat WebFilter database, and the ProxySG expediently updates its on-box copy of the Symantec WebFilter database. About 95% of the Web requests made by a typical enterprise user (for the English language) are present in the on-box Symantec WebFilter database, thereby minimizing bandwidth usage and maintaining quick response times.

By default, the WebPulse service is enabled and configured to dynamically categorize unrated and new Web content for immediate enforcement of policy. Typically, the response time from the dynamic categorization service is about 500 milliseconds and is subject to the response/performance of the site in question. However, if this service is causing significant delays to your enterprise Web communications, you can run it in Background mode.

If you disable dynamic categorization, proactive threat detection, content and reputation ratings are also disabled.

Note: Symantec WebFilter customers have the option to use a secure connection for establishing communication between the dynamic categorization client on the ProxySG and the WebPulse service. For information, see "[Configuring WebPulse Services](#)" on page 401.

By default, connection to the WebPulse service is not encrypted and data is sent as plain text; however, you can opt to use a secure connection, which encrypts all data sent over the connection.

About Dynamic Categorization

The dynamic categorization service analyzes and categorizes new or previously unknown URLs, which are not in the on-box Symantec WebFilter database. Dynamic categorization can be processed in two modes—immediately or in the background.

By default, dynamic categorization is set to be performed immediately, which is in real time. When a user requests a URL that has not already been categorized by the Symantec WebFilter database (for example, a new Web site), the WebPulse dynamic categorization service queries the target Web site and retrieves the page's content. WebPulse analyzes the page's content and context in search of malicious content. If malicious content is found, an appropriate category (for example, Spyware/Malware sources or Phishing) is returned. If no malicious content is found, WebPulse's dynamic real time rating (WebPulse) service

determines the language of the page, a category for the page, and a confidence factor that the category is correct. If the confidence factor is high, the calculated category is returned.

In situations where the dynamic categorization service cannot categorize a URL with enough confidence to dynamically return a category with a high confidence level, the category rating request for the particular page is labeled *none*. All URLs received by WebPulse that are not categorized in the Symantec WebFilter database are logged and forwarded to Symantec's centralized processing center, where they are prioritized for rating by a series of automated URL analysis tools and/or human analysis. These ratings are then used to update the master Symantec WebFilter database, and the automatic database update feature then refreshes the local Symantec WebFilter database on the ProxySG.

When dynamic categorization is performed in Background mode, the ProxySG continues to service the URL request without waiting for a response from the WebPulse dynamic categorization service. The system category *pending* is assigned to the request, indicating that the policy was evaluated with potentially incomplete category information. When WebPulse returns a category rating, the rating is stored in a dynamic categorization cache so that the next time the URL is accessed, WebPulse will not be required to determine its category.

Note: The dynamic service is consulted *only* when the installed Symantec WebFilter database does not contain authoritative category information for a requested URL. If the category returned by the WebPulse service is blocked by policy, the offending material does not re-enter the network.

The URL is first looked up in the local Symantec WebFilter database. The expected results are shown in the following table.

<i>Found in Local WebFilter Database</i>	<i>Found in Rating Cache</i>	<i>Process Mode</i>	<i>Result / Description</i>
Yes	Any	Any	The corresponding list of categories in the database is returned.
No	Yes	Any	The corresponding list of categories in the ratings cache is returned.
No	No	Dynamic Categorization disabled	None. No categories are available for the URL.

<i>Found in Local WebFilter Database</i>	<i>Found in Rating Cache</i>	<i>Process Mode</i>	<i>Result / Description</i>
No	No	Dynamic Categorization in Background Mode	Pending. The ProxySG appliance continues to service the URL request without waiting for a response from WebPulse. If a response is received, it is added to the rating cache, so future requests for that same URL will have the appropriate list of categories returned immediately. Reference to the site is recorded for future categorization in the WebFilter database by automated background URL analysis or human analysis. If a response is not received in a timely manner, or the request results cannot be categorized, nothing is added to the rating cache. Note: It is possible that multiple requests for the same content can result in a Pending status if WebPulse has not completed processing the first request before subsequent requests for the same URL are received by the ProxySG appliance.
No	No	Dynamic Categorization in Real-time Mode and categories returned by WebPulse	A request to categorize the URL is sent to WebPulse and the ProxySG appliance waits for a response. The response is added to the rating cache and also used as the list of categories for the current request.
No	No	Dynamic Categorization in Real-time Mode and categories not returned by WebPulse	None. Categories might not be returned because: • The ProxySG appliance did not get a response from the WebPulse service. • The WebPulse service was unable to retrieve the requested URL in a timely manner. • The WebPulse service cannot categorize the request with high confidence. References to all URLs requested in WebPulse are recorded for future categorization in WebFilter by automated background analysis or human analysis. Note: Timeout is currently set to three seconds. Average response time for WebPulse to retrieve the content and perform real-time analysis is under 500 milliseconds.

<i>Found in Local WebFilter Database</i>	<i>Found in Rating Cache</i>	<i>Process Mode</i>	<i>Result / Description</i>
Any	Any	Any	Unlicensed. A problem exists with the WebFilter license.
Any	Any	Any	Unavailable. A problem (other than licensing) exists with the local WebFilter database or accessing the WebPulse service.

See Also:

- ❑ ["About the Dynamic Categorization Process"](#) on page 385
- ❑ ["Dynamic Categorization States"](#) on page 386
- ❑ ["Considerations Before Configuring WebPulse Services"](#) on page 387
- ❑ ["About Private Information Sent to WebPulse"](#) on page 388

About the Dynamic Categorization Process

Dynamic analysis of content is performed through the WebPulse cloud service and not locally on the ProxySG. There is a small amount of bandwidth used for the round-trip request and response, and a slight amount of time waiting for the service to provide results. As the service is only consulted for URLs that cannot be locally categorized using the Symantec WebFilter database and WebPulse results are cached on the ProxySG, the user experience is generally not affected.

To avoid per-request latency, you might want to run dynamic categorization in *background mode*. For modifying the default, see ["Configuring WebPulse Services"](#) on page 401.

The following diagram illustrates Symantec WebFilter's content filtering flow when dynamic categorization is employed.




Figure 20–2 WebFilter with Dynamic Categorization Content Enabled (default)

Dynamic Categorization States

Dynamic categorization has three states:

- ❑ **Enabled:** The service attempts to categorize unrated Web sites. This is the default state.
- ❑ **Disabled:** If the service is disabled, the ProxySG does not contact the WebPulse service, regardless of any policy that might be installed.

- ❑ **Suspended:** Categorization from the database continues, but the service is no longer employed. This occurs when the installed database is over 30 days old due to the expiration of WebFilter download credentials or network problems. After credentials are renewed or network problems are resolved, the service returns to Enabled.

Considerations Before Configuring WebPulse Services

The WebPulse protocol regulates the communication between the dynamic categorization client on the ProxySG and the WebPulse cloud service. Before configuring WebPulse services using ["Configuring WebPulse Services"](#) on page 401, answer the following questions:

- ❑ Do you use proxy chaining or SOCKS gateways?
If you use a forwarding host for forwarding dynamic categorization requests through upstream proxies or use SOCKS gateways, see ["About Proxy Chaining Support for WebPulse Services"](#) for more information on the forwarding options in WebPulse. For information on forwarding and configuring the upstream network environment, see ["Configuring the Upstream Network Environment"](#) on page 865.
- ❑ Would you like to configure private networks to identify traffic relating to your internal networks while using WebPulse for content rating accuracy?
If you specify your private networks, information pertaining to the configured internal network is removed by the ProxySG appliance, prior to sending a dynamic categorization request across to the WebPulse service. For understanding the interaction between private networks and dynamic categorization, see ["About Private Information Sent to WebPulse"](#) on page 388. To configure private networks for maintaining your security needs, see ["Configuring Private Networks"](#) on page 825.
- ❑ Would you like to provide malware feedback notification to the WebPulse community?
This option is applicable only if you have a Blue Coat AV configured on the ProxySG, and malware scanning and WebFilter are enabled. When the ProxySG is integrated with the Blue Coat AV, the ProxySG monitors the results of the Blue Coat AV scan and notifies the WebPulse service when a new virus or malware is found. This feedback helps update the malware and content ratings and protects the entire community of WebPulse users. For more information, see ["About Malware Notifications to WebPulse"](#) on page 391.
For information on adding a Blue Coat AV and enabling malware scanning, see ["Adding an ICAP service for Content Scanning"](#) and ["Enabling Malware Scanning"](#)

About Proxy Chaining Support for WebPulse Services

Proxy chaining is a method for routing client requests through a chain of ProxySG appliances until the requested information is either found in cache or is serviced by the OCS.

The ProxySG allows you to forward dynamic categorization requests through upstream proxies and SOCKS gateways.

Important: When configuring a forwarding host under **Configuration > Forwarding > Forwarding Hosts**, in the **Add Forwarding Host** dialog select **Type: Proxy**. If you attempt to configure proxy chaining using **Type** as **Server**, an error occurs.

Forwarding Hosts and Dynamic Categorization

To forward dynamic categorization requests through an upstream HTTP proxy, configure a forwarding host that is defined as a proxy and specify the HTTP port for the connection. You can then select that forwarding host in the WebPulse configuration.

Note: If forwarding is configured, you cannot enable *secure* dynamic categorization; if secure dynamic categorization is enabled, you cannot select a forwarding host.

SOCKS Gateways

If you use proxy chaining for load balancing or for forwarding the dynamic categorization request through an upstream SOCKS gateway, you must configure the SOCKS gateway before configuring the WebPulse service.

Important: Before configuring the SOCKS gateway target for WebPulse, verify that the SOCKS gateway is operating correctly.

When both SOCKS and forwarding are configured, the ProxySG connects to the SOCKS gateway first, then to the forwarding host, and then to the WebPulse service.

About Private Information Sent to WebPulse

A private network is an internal network that uses private subnets and domains, for example, your intranet. On the ProxySG, you can configure private networks on the **Configuration > Network > Private Network** tab. For information on configuring private subnets or private domains, see "[Configuring Private Networks](#)" on page 825.

By default, dynamic categorization is enabled on the ProxySG. When a requested URL is not in the dynamic categorization ratings cache or categorized in the WebFilter database, the request is sent to the WebPulse cloud service for dynamic categorization. By configuring private subnets and private domains within your network, you can use dynamic categorization to ensure accuracy of content filter ratings, while preserving the security of sensitive information relating to your private networks.

Before a request is sent for content rating to the WebPulse cloud service, the following conditions are verified on the ProxySG appliance:

- Is WebPulse service and dynamic categorization enabled?
- Is dynamic categorization permitted by policy?
- Is the host specified in the private domain or private subnet list?

Any request that is determined to be part of your configured private network is not sent to WebPulse.

The ProxySG appliance may send information from HTTP and HTTPS requests to the WebPulse service if they are not directed to hosts that are part of the configured private network.

Note: Private network domain names and IP subnets can be user-defined.

Customer information sent to the WebPulse service is controlled by user-defined policy, although you can still use the default policy and configuration settings provided by the ProxySG appliance. Overriding the default settings with your organization's policy definitions results in more control of the type of information that is sent to the WebPulse service.

When WebPulse service is enabled, the default configuration settings in SGOS 5.4.1 and later send the fixed customer information listed above, in addition to the following information:

- URL query string
- Referer header
- User-Agent header

However, this additional information can be controlled by policy and/or configuration settings:

If the `service send-request-info` setting is set to `disable`, by default, only the customer license key, URL scheme, method, host, port, and path are sent to the WebPulse service; URL query string, and the `Referer` and `User-Agent` headers are not sent.

If the `service send-request-info` setting is set to `enable`, by default, referrer and user agent information and so on can be sent to the server. Symantec gathers this customer information from back-end logs and analyzes the data to help improve its threat protection. In its analysis, Symantec does not consider the source of the data; that is, customer information is anonymous.

Note: Be aware that personal information might be included in the URL query string. If this information is sent to WebPulse, Symantec might use it when accessing content from the Web site in order to categorize it.

You can further control whether to include the URL path and query string, and individually control whether the `Referer` or `User-Agent` headers are sent for specific requests. Restrictions are accomplished through the use of policies that can be defined from the ProxySG appliance management console or CLI.

Table 20–1 on page 390 lists the type of information that is sent to the WebPulse service based on default settings for all SGOS versions supporting WebPulse.

Note: The `service send-request-info` command applies only to SGOS 5.4.1 and later.

Table 20–1 Information Sent to the WebPulse Service Based on Default SGOS Settings

Information Sent to the WebPulse Service	SGOS < 5.4.1	SGOS 5.4.1 and Later (service send-request-info disable)	SGOS 5.4.1 and Later (service send-request-info enable)	SGOS 6.4 and Later (service send-request-info disable)	SGOS 6.4 and Later (service send-request-info enable)
Customer License Key (Example: QA852-KL3RA)	Yes	Yes	Yes	Yes	Yes
Scheme (Examples: HTTP, HTTPS)	Yes	Yes	Yes	Yes	Yes
Method (Examples: GET, POST)	Yes	Yes	Yes	Yes	Yes
URL Host/Port	Yes	Yes	Yes	Yes	Yes
URL Path ¹	Yes	Yes	Yes ²	Yes	Yes ²
URL Query String	No	No	Yes ²	No	Yes ²
Referer Header	No	No	Yes ²	No	Yes ²
User-Agent Header	No	No	Yes ²	No	Yes ²
Content-Type	No	No	No	No	Yes ²
Content-Length	No	No	No	No	Yes ²

1. Path = URL minus any query string.

2. Can be controlled using Content Policy Language (CPL).

See Also

- ❑ "Configuring WebPulse Services" on page 401
- ❑ "Viewing Dynamic Categorization Status (CLI only)" on page 404

- ❑ [Section H: "Applying Policy"](#) on page 414
- ❑ [Content Policy Language Reference](#)

About Malware Notifications to WebPulse

The Blue Coat AV, when integrated with the ProxySG, provides in-path threat detection. The Blue Coat AV scans Web content based on the protection level in your malware scanning configuration in **Configuration > Threat Protection> Malware Scanning**. Every proxied transaction is scanned when the protection level is set at maximum security; selected transactions are subject to a monitoring check when the protection level is set to high performance.

By default, if a malware threat is detected, the Blue Coat AV notifies the ProxySG, which then issues a malware notification to the WebPulse service. This notification triggers an update of the WebFilter database, and all members of the WebPulse community are protected from the emerging threat. When malware scanning is enabled, notification requests sent to the WebPulse service include the request URL, HTTP `Referer` and `User-Agent` headers.

If the `service send-malware-info` setting is set to `enable` (by default), the customer information listed above is sent to the WebPulse service by the ProxySG appliance.

If the `service send-malware-info` setting is set to `disable`, malware notification requests are not sent to the WebPulse service by the ProxySG appliance.

Note: Symantec respects your security needs. If the request URL or the `Referer` header for a malware threat pertains to a private URL, no malware notification is issued.

See Also

- ❑ ["Configuring WebPulse Services"](#) on page 401
- ❑ ["Viewing Dynamic Categorization Status \(CLI only\)"](#) on page 404

Section B: Setting up a Web Content Filter

This section provides the list of tasks required to configure a Web content filter for monitoring, managing, and restricting access to Web content. The following topics are discussed:

- ❑ ["Web Content Filtering Task Overview"](#) on page 392
- ❑ ["Enabling a Content Filter Provider"](#) on page 392
- ❑ ["Downloading the Content Filter Database"](#) on page 394
- ❑ ["Setting the Memory Allocation"](#) on page 398

Web Content Filtering Task Overview

Before you begin setting up content filtering, ensure that you have a valid subscription from a content filter provider of your choice. Only the IWF and local databases do not require a subscription or license.

To set up on-box Web content filtering on the ProxySG, perform the following tasks:

- ❑ ["Enabling a Content Filter Provider"](#)
- ❑ ["Downloading the Content Filter Database"](#)
- ❑ ["Applying Policy"](#)

To review the default settings for your content filtering vendor and to make adjustments, see the following sections:

- ❑ ["Configuring Symantec WebFilter and WebPulse"](#): The default settings are adequate for most environments. This section provides information on customizing WebFilter settings to meet the needs in your network.
- ❑ ["Configuring a Local Database"](#): A local database is typically used in conjunction with a standard content filter database that has pre-defined categories. This section provides information on creating and maintaining a local database for your network.
- ❑ ["Configuring Internet Watch Foundation"](#): This section provides information on customizing the download schedule for the IWF database that includes a single category called **IWF-Restricted**.


Enabling a Content Filter Provider

This is the first step in setting up a Web content filter on the ProxySG. This procedure assumes you have a valid account with your preferred vendor.

Prerequisite: ["Web Content Filtering Task Overview"](#) on page 392

To enable a content filter provider:

1. Select the **Configuration > Content Filtering > General** tab.



2. Select the option for your preferred provider. You can opt to enable the local database, Internet Watch Foundation, WebFilter, a third-party vendor (select your preferred vendor from the **Third-party database** drop-down list), and YouTube.

Note: You cannot configure the legacy content filtering vendors, such as Intersafe, I-Filter, and Surfcontrol, using the Management Console. If you upgraded the ProxySG appliance from an SGOS version previous to 5.5.x, and one of the vendors listed above is enabled on your ProxySG, use the Command Line Interface (CLI) to modify configuration settings for these vendors.

3. Select the **Lookup Mode** option. For a Web request, the look up mode determines the databases that the ProxySG searches for a category match. To perform a lookup, the database must be enabled. The look up sequence executed is policy, local database, IWF, Symantec WebFilter and finally a selected third-party database.

Note: For YouTube, the Lookup mode option is hard-coded to **Always**. This means that the database is always consulted for category information.

- a. The default is **Always**, which specifies that the database is always consulted for category information. If a URL is categorized under more than one category in different databases, policy is checked against each category listed.
- b. **Uncategorized** specifies that a database lookup be skipped if the URL match is found in policy, a Local database, or the Internet Watch Foundation (IWF) database.

4. (Applicable for WebFilter *only*) Select **Enable category review message in exceptions**. This option adds a link to the default content filter exception page when a user is denied a request for a Web Page. Typically the exception page informs the user why a URL request is denied. When you enable this option, the user can click the link displayed on the exception page to request a review of the category assigned to the blocked URL. For example, when enabled the screen displays the following users:

Your request was categorized by Blue Coat WebFilter as 'News/Media'.
If you wish to question or dispute this result, please click [here](#).

The built in exception page can be customized, for customizing the exception page, refer to the Advanced Policy Tasks chapter, Section E, of the *Visual Policy Manager Reference*.

5. Click **Apply**.

Downloading the Content Filter Database

The dynamic nature of the Internet makes it impossible to categorize Web content in a static database. With the constant flow of new URLs, URLs of lesser-known sites and updated Web content, maintaining a current database presents a challenge. To counter this challenge, the ProxySG supports frequent content filter database downloads.

For more information, see one of the following topics:

- ["About Database Updates"](#)
- ["Downloading a Content Filter Database"](#) on page 395

For more information about the Symantec WebFilter database, see ["About Symantec WebFilter and the WebPulse Service"](#) on page 381.

About Database Updates

Symantec enables all customers with a valid content filtering license to schedule automatic downloads of content filter databases. By default, automatic updates are enabled; The ProxySG appliance checks for updates once in every five minutes and downloads an incremental update when available.

After selecting your provider(s) of choice, you must enter the license information and download the database(s) on the ProxySG. You can download a database on demand or schedule a periodic download using the automatic download feature.

Typically, a complete database download occurs when you enable the provider and add the license key for the first time. Thereafter, the ProxySG periodically checks the download server for updates to the installed database. If the database is current, no download is performed.

When an update is available, it is automatically downloaded and applied. An update provides the most current categorization of URLs and contains only the changes between the current installed version and the latest published version of the database, and hence is much smaller than a full copy of the database. In the unlikely event that this conditional download fails, the ProxySG downloads the latest published version of the complete database.

Note: By default, the ProxySG checks for database updates once in every five minutes. While you can schedule the time interval for an automatic database update, the frequency of checks is not configurable.

Continue with ["Downloading a Content Filter Database"](#).

Downloading a Content Filter Database

This section discusses how to download the following content filter databases through the Management Console:

- ❑ Symantec WebFilter
- ❑ Internet Watch Foundation (IWF)
- ❑ Proventia
- ❑ Optenet

Note: To download the Surfcontrol, I-Filter, or Intersafe databases, use the Command Line Interface (CLI). Refer to the *Command Line Interface Reference* for a list of commands.

For information about content filter updates, or if you are setting up the content filter provider for the first time, see ["About Database Updates"](#) on page 394.

To download the content filter database:

1. If you are downloading the WebFilter or IWF database, select the **Configuration > Content Filtering > Vendor_Name** tab. Alternatively, if you are downloading a third-party vendor database, select the **Configuration > Content Filtering > Third-Party Databases > Vendor_Name** tab (this example uses Optenet).

The screenshot shows the Management Console interface for downloading a content filter database. The 'Optenet' tab is active. The 'Download' section contains the following elements:

- Username:** A text input field.
- Change Password:** A button.
- URL:** A text input field containing the URL `https://list.bluecoat.com/optenet/activity/download/optenet.db`.
- Set to default:** A button.
- Download now:** A button.
- View Download Status:** A button.
- Automatically check for updates:** A checked checkbox.
- Only between the hours of:** A checkbox followed by two time selection dropdowns set to '00:00' and '23:59'.

2. Download the database. Except for IWF, you must enter valid subscription credentials to download the database. If the database has previously been downloaded on a local Web server that requires authentication, you must configure the ProxySG to use credentials that allow access to the Web server, which hosts the content filter database.
 - a. Enter your username and password (required for WebFilter, Proventia, and Optenet).
 - b. (Optional) Click **Change Password**. The Change Password dialog displays. Enter your password and click **OK**.
 - c. The default database download location is displayed in the **URL** or **Server** field. If you have been instructed to use a different URL, enter it here.
 - d. (Optional) If you changed the URL for downloading the database, to reset to the default location, click **Set to default**. The default download location overwrites your modification.
 - e. Click **Apply** to save all your changes.
 - f. Click **Download Now**. The **Download Status** dialog displays.
 - g. Click **Close** to close the Download status dialog.

It may take several minutes for the database download to complete. When the database has been downloaded, proceed to ["Viewing the Status of a Database Download"](#).

Viewing the Status of a Database Download

When the database is downloaded, the download log includes detailed information on the database.

If you have just configured content filtering and are downloading the database for the first time, the ProxySG downloads the latest published version of the complete database. Subsequent database updates occur incrementally.

To view the status of the download:

On the **Configuration > Content Filter > Vendor_Name** tab, click **View Download Status**. A new browser window opens and displays the download log. For example:

```
Download log:
  Optenet download at: 2009/09/11 17:25:52 +0000
  Downloading from https://list.bluecoat.com/optenet/activity/
  download/optenet.db
  Warning: Unable to determine current database version; requesting
  full update
  Download size:      37032092
  Database date:      Thu, 10 Sep 2009 09:30:44 UTC
  Database expires:   Sat, 10 Oct 2009 09:30:44 UTC
  Database version:   1629
  Database format:    1.1
```

Expiry Date for the Database

A valid vendor subscription is required for updating your database. Each time a database download is triggered manually or using the automatic download feature, the validity of the database is reset.

When your license with the database vendor expires, you can no longer download the latest version. The expiry of a database license does not have an immediate effect on performing category lookups for the on-box categories. You can continue to use the on-box database until the expiry of the database. However, when the database expires, the category **unlicensed** is assigned to all URLs and no lookups occur on the database.

Viewing the Available Categories or Testing the Category for a URL

For each content filter vendor whose database has been downloaded on the ProxySG, you can view the list of categories available. This list is relevant for creating policy that allows or restricts access to Web content and for verifying the category that a URL matches against in the database.

To view the available categories for a content filter vendor:

1. Select the **Configuration > Content Filtering > General** tab.
2. Click **View Categories**. The list of categories displays in a new Web page.

To verify the category assigned to a URL:

1. Select the **Configuration > Content Filtering > General** tab.
2. Enter the URL into **URL**.
3. Click **Test**. A new Web page displays with the category that your chosen vendor(s) has assigned to the URL. For example, the URL `cnn.com` is categorized as follows:
Blue Coat: News/Media
Optenet: Press

Note: The maximum number of categories for any single URL is 16. If more than 16 categories are specified, the ProxySG appliance arbitrarily matches against 16 out of the total number specified.

Testing the Application and Operation for a URL

If you are using WebFilter for content filtering, you have the additional ability to deny or allow access to certain web applications and/or operations; this is done via policy—either using the VPM or CPL. If you want to find out the application or operation name associated with a URL so that you can create policy to block or allow it, you can do a URL test, as described below.

To determine the category, application, and operation associated with a URL:

1. Select the **Configuration > Content Filtering > General** tab.
2. Enter the URL into **URL**.

3. Click **Test**. A new Web page displays with the category, application, and operation that WebFilter has assigned to the URL. For example, the URL `facebook.com/video/upload_giver.php` is categorized as follows:
Social Networking; Audio/Video Clips
Facebook
Upload Videos

The test results in this example indicate that the URL has two categories (Social Networking and Audio/Video Clips), is the Facebook application, and is the Upload Videos operation.

Note that not all URLs have applications and operations associated with them. For URLs that WebFilter has not assigned an application or operation, the test results indicate `none`.

Setting the Memory Allocation

Note: The default memory allocation (normal) setting is ideal for most deployments. This procedure is relevant only to specific deployments as detailed below.

Content filtering databases can be very large and require significant resources to process. It might be necessary to adjust the amount of memory allocated to the database in the following situations:

- ❑ If you are *not* using ADN and have a high transaction rate for content filtering, you can increase the memory allocation setting to High. This helps content filtering run more efficiently.
- ❑ If you are using both ADN and content filtering but the transaction rate for content filtering isn't very high, you can reduce the memory allocation setting to Low. This makes more resources available for ADN, allowing it to support a larger number of concurrent connections.

To set the memory allocation for content filtering:

1. Select the **Configuration > Content Filtering > General** tab.
2. Select the memory allocation setting that works for your deployment: **Low**, **Normal**, or **High**.
3. Click **Apply**.

Section C: Configuring Symantec WebFilter and WebPulse

This section describes how to modify the defaults for Symantec WebFilter, customize your database update schedule, and modify the WebPulse service that detects malware threats and controls real-time rating of client requests.

Important: WebFilter requires a valid license. For information on Licensing, see ["Licensing"](#).

Configuring Symantec WebFilter

Symantec WebFilter is an on-box content filtering database that protects data and users from network attacks. All Symantec WebFilter subscribers are a part of the WebPulse cloud service, which continuously updates the on-box database. Symantec WebFilter and WebPulse provide Dynamic Real-Time Rating, a technology that can instantly categorize Web sites when a user attempts to access them.

The following sections describe making adjustments to the Symantec WebFilter defaults:

- ❑ ["Disabling Dynamic Categorization"](#) on page 399
- ❑ ["Specifying a Custom Time Period to Update Symantec WebFilter"](#) on page 400
- ❑ ["Configuring WebPulse Services"](#) on page 401
- ❑ ["Viewing Dynamic Categorization Status \(CLI only\)"](#) on page 404

See Also


- ❑ ["About Dynamic Categorization"](#) on page 382
- ❑ ["About Private Information Sent to WebPulse"](#) on page 388

Disabling Dynamic Categorization


By default, when you enable and download the Symantec WebFilter database, dynamic categorization in real time is available on the ProxySG.

To disable dynamic categorization:

1. Select the **Configuration > Content Filtering > Blue Coat WebFilter** tab. For information on enabling and downloading Symantec WebFilter, see ["Enabling a Content Filter Provider"](#) and ["Downloading the Content Filter Database"](#).



2. Click the **Dynamic categorization** link. The **Configuration > Threat Protection > WebPulse** tab displays.



3. Clear the **Perform Dynamic Categorization** option in the **Configuration > Threat Protection > WebPulse** page. If you disable dynamic categorization, proactive threat detection, content and reputation ratings are also disabled. For information on dynamic categorization, see "[About Dynamic Categorization](#)" on page 382. For information on performing dynamic categorization in background mode, see "[Configuring WebPulse Services](#)"

Specifying a Custom Time Period to Update Symantec WebFilter

Database updates provide you with the most comprehensive and current URL categories. The ProxySG checks for database updates in five minute intervals.

The automatic download setting is enabled by default, but you can disable this feature if desired. You can customize the window of time at which the automatic update happens, for example, you might specify automatic updates only between the hours of 8 pm and 11 pm. The time frame is always local time. Note that the frequency of updates within the specified time period is *not* configurable.

To specify a custom time period for updates:

1. Select the **Configuration > Content Filtering > Blue Coat** tab. The **Automatically Check for Updates** option is selected by default.

Blue Coat WebFilter

Dynamic categorization: [Enabled](#)

Download

Username:

Change the download password

URL:

☒ Automatically check for updates

☒ Only between the hours of and

2. Configure the options:
 - a. Select the **Only between the hours of** option. The time frame is local time.
 - b. Expand the drop-down lists, and set the time period for your update schedule. For example, to check for updates between the hours of 7 pm and midnight, set the first box to **19:00** and the second box to **23:59**.
3. Click **Apply**.

Note: The update check frequency configuration is an available setting in each of the supported content filter providers.

Configuring WebPulse Services

WebPulse is a cloud service that provides the off-box component of Symantec's complete content filtering solution. The WebPulse cloud service blocks malware hosts, rates Web content and protects both ProxySG appliance Web gateways and remote users of ProxyClient. For more information, see "[About Symantec WebFilter and the WebPulse Service](#)".

This section describes how you can modify or disable dynamic categorization settings, enable secure connections for WebPulse services, and disable the malware feedback loop between the Blue Coat AV and the ProxySG.

To configure WebPulse services:

1. Select the **Configuration > Threat Protection > WebPulse** tab.

The screenshot shows the WebPulse configuration page. On the left, there are numbered callouts (2 through 7) with arrows pointing to specific settings:

- 2** points to the ☒ **Enable WebPulse service** checkbox.
- 3** points to the **Blue Coat WebFilter: Enabled** text and the **Last download: Unknown** text.
- 4** points to the ☐ **Use secure connections** checkbox.
- 5** points to the **Forwarding target:** dropdown menu, which currently shows **none**.
- 6** points to the **Dynamic Categorization** section, which includes:
 - ☒ **Perform dynamic categorization**
 - ☐ **Immediately** (Categorize the request in real-time and wait for the result.)
 - ☒ **In the background** (Categorize the request but do not wait for the result, returning the category as "pending" on the initial access.)
- 7** points to the ☒ **Send potential malware sources to WebPulse** checkbox in the **Malware Feedback** section.

2. (Optional) To disable the WebPulse service, clear **Enable WebPulse Service**. If you disable the WebPulse service, dynamic categorization and malware feedback are also disabled. References to perform dynamic categorization in policy are also disregarded. For information on the WebPulse service, see ["About Symantec WebFilter and the WebPulse Service"](#) on page 381.
3. Verify that Symantec WebFilter is enabled as your content filter vendor and confirm that the last download was completed within the previous 24-hour interval.
4. (Optional) To enable secure connections (data encryption), select **Use secure connections**. If you use a secure connection, you cannot select a forwarding host or group.
You can, however, use a SOCKS gateway for establishing secure connections to the WebPulse cloud service. For information on proxy chaining, see ["About Proxy Chaining Support for WebPulse Services"](#) on page 387.

Note: For most situations, using secure connections does not significantly decrease performance unless you are regularly processing a large number of unrated sites.

5. (Optional) Select a forwarding host or a SOCKS gateway target. You cannot select a forwarding host or group if you enabled secure connections in Step 4.

6. To modify the dynamic categorization mode, verify that the **Perform Dynamic Categorization** option is selected and Symantec WebFilter is enabled. Then choose one of the following options:

- a. **Immediately.** This is the default categorization mode and is in real-time — if the category of the request is not already known, the URL request will wait for the WebPulse service to respond with the categorization before proceeding. The advantage of real-time mode categorization is that Blue Coat policy has access to the results, allowing policy decisions to be made immediately after receiving all available information.
- b. **In the background.** In this mode when dynamic categorization is triggered, the URL request continues to be serviced without waiting for a response from the WebPulse service. The system category *pending* is assigned to the request, indicating that the policy was evaluated with potentially incomplete category information.

The result of the categorization response is entered into a categorization cache; This cache ensures that any subsequent requests for the same or similar URLs can be categorized quickly, without needing to query the WebPulse cloud service again.

Note: If Symantec WebFilter license has expired and dynamic categorization is enabled, the service enters a suspended state. For more information, see "[Dynamic Categorization States](#)".

- c. (Optional) To disable dynamic categorization, clear the **Perform Dynamic Categorization** checkbox.

If dynamic categorization is disabled, the ProxySG does not contact the WebPulse service when a category match for a URL is not found in the on-box database.

- d. (Optional) Disable **Malware Feedback**. This setting is enabled by default for all dynamic categorization requests. If you have a ProxySG integrated with the Blue Coat AV for ICAP scanning and WebFilter and WebPulse are enabled, when a malware threat is detected the Blue Coat AV notifies the ProxySG. The ProxySG then issues a malware notification to WebPulse for updating the Symantec WebFilter database.

When this option is disabled, the ProxySG does not notify WebPulse about malware URLs that the Blue Coat AV detects. However, you can use policy to override the default malware feedback settings.

7. Click **Apply**.

Configuring Dynamic Categorization Requests for HTTP/HTTPS (CLI only)

You can configure dynamic categorization requests for HTTP and HTTPS transactions sent to WebPulse in the CLI.

To enable or disable sending HTTP header information to WebPulse, use the following command:

```
SGOS#(config bluecoat) service send-request-info {enable | disable}
```

The setting is enabled by default. When enabled, WebPulse receives HTTP headers with `Referer`, `User-Agent`, `Content-Type`, and `Content-Length` information. When the setting is disabled, the ProxySG does not send any HTTP header information to WebPulse.

To specify the mode and amount of information sent to WebPulse for HTTPS transactions, use the following command:

```
SGOS#(config bluecoat) service send-https-url {full | path | disable}
```

The following are parameters for the command:

- `full` — Send entire URL (domain, path, and query string).
- `path` — Send only the domain and path.
- `disable` — Do not send a rating request for HTTPS transactions.

Viewing Dynamic Categorization Status (CLI only)

The dynamic categorization feature has three states—enabled, disabled, and suspended.

When enabled, the ProxySG accesses the WebPulse cloud service for categorizing a requested URL when it is not available in the Symantec WebFilter database.

When disabled or suspended, the ProxySG does not access the WebPulse cloud service for categorizing a requested URL. The Symantec WebFilter database is consulted for categorization and based on the policies installed on the ProxySG, the requested content is served or denied.

Service suspension occurs when the installed database is over 30 days old. The main reasons for service suspension are the expiration of Symantec WebFilter download credentials or due to network problems in downloading the latest database version. When the credentials are renewed or network problems are resolved, the service returns to Enabled.

To view the dynamic categorization status, at the `(config)` prompt, enter the following command:

```
SGOS# (config content-filter) view  
Provider: Blue Coat  
Dynamic Categorization:  
Service: Enabled/Disabled/Suspended <---one state is displayed
```

See Also

- ❑ ["Applying Policy"](#)
- ❑ ["Applying Policy to Categorized URLs"](#)
- ❑ ["More Policy Examples"](#)
- ❑ ["Defining Custom Categories in Policy"](#)

Section D: Configuring a Local Database

The following sections describe how to select and refer to a local database and how to schedule the database update schedule:

- ❑ ["About the Local Database"](#)
- ❑ ["Local Database Matching Example"](#)
- ❑ ["Selecting and Downloading the Local Database"](#)

About the Local Database

Two main reasons to use a local database instead of a policy file for defining categories are:

- ❑ A local database is more efficient than policy if you have a large number of URLs.
- ❑ A local database separates administration of categories from policy. This separation is useful for three reasons:
 - It allows different individuals or groups to be responsible for administering the local database and policy.
 - It keeps the policy file from getting cluttered.
 - It allows the local database to share categories across multiple boxes that have different policy.

However, some restrictions apply to a local database that do not apply to policy definitions:

- ❑ No more than 200 separate categories are allowed.
- ❑ Category names must be 32 characters or less.
- ❑ A given URL pattern can appear in no more than four category definitions.
- ❑ The local database produces only the most specific URL match and returns a single category.

The same policy syntax will produce a different match. If more than one category is provided, policy processing may match more than one category and hence will return more than one category. See ["Local Database Matching Example"](#) on page 406 for more information.

You can use any combination of the local database, policy files, or the VPM to manage your category definitions. See ["Applying Policy to Categorized URLs"](#) on page 414 for more information. You can also use both a local database and a third-party vendor for your content filtering needs.

Note: Blue Coat recommends locating your local database on the same server as any policy files you are using.

Local Database Matching Example

As noted above, the local database produces only the most specific URL match and returns a single category. Consider the following examples.

Local Database Example

Consider the following syntax.

```
define category no_detect_protocol
mail.google.com
end
```

```
define category google
google.com
end
```

Local database result:

```
https://<proxy>:8082/ContentFilter/TestUrl/mail.google.com/
  Local: no_detect_protocol
  Blue Coat: Mail
Gmail
none
```

Policy Example

This example uses the same syntax as the local database example.

```
<proxy>
ALLOW
```

```
define category no_detect_protocol
mail.google.com
end
```

```
define category google
google.com
end
```

Policy Result:

```
https://<proxy>:8082/ContentFilter/TestUrl/mail.google.com/
  Policy: no_detect_protocol; google
  Blue Coat: Mail, Search Engine
Gmail
none
```

As shown, policy returns both categories; whereas, the local database returns only the URL match.

Creating a Local Database

The local database is a text file that must be located on a Web server that is accessible by the ProxySG appliance on which you want it configured. You cannot upload the local database from a local file.

The local database file allows `define category` statements only.

To create a local database:

1. Create a text file in the following format:

```
define category <category-name>
url1
url2
urln
end

define category <category-name>
url1
url2
urln
end
```

Each category can have an unlimited number of URLs.

For example,

```
define category bluecoat_allowed
bluecoat.com
symantec.com
yahoo.com
microsoft.com
sophos.com
end

define category bluecoat_denied
www.playboy.com
www.hacking.com
www.sex.com
www.poker.com
'[2607:F330:8500:220::195]'
'216.139.0.95'
end
```

2. Upload the text file to a Web server that the ProxySG appliance can access.
3. Continue with ["Selecting and Downloading the Local Database"](#).

Selecting and Downloading the Local Database

This section discusses how to select a local database to serve your content filtering needs. To create the local database, see ["Local Database Matching Example"](#) on page 406.

To configure local database content filtering:

1. Select the **Configuration > Content Filtering > General** tab.

Provider	Enable	Lookup mode
Local Database:	☒	☒ Always ☐ Uncategorized
Internet Watch Foundation:	☐	☒ Always ☐ Uncategorized
Blue Coat Web Filter:	☐	☒ Always ☐ Uncategorized
3rd-party database:	None	☒ Always ☐ Uncategorized

Options

☐ Enable Category Review Message in Exceptions

2. Select **Local Database**.
3. Select the **Lookup Mode**:
 - a. The default is **Always**, which specifies that the Local database is always consulted for category information.
 - b. **Uncategorized** specifies that the lookup is skipped if the URL has already been found in policy.
4. Click **Apply**.
5. Select the **Configuration > Content Filtering > Local Database** tab.
6. If the database is located on a server that requires a password for access, you must configure the ProxySG to use that password when accessing the database:
 - a. Click **Change Password**. The Change Password dialog displays.
 - b. Enter your password and click **OK**.
7. Download the database:
 - a. In the **URL** field, enter the location of the file to be downloaded.
 - b. Click **Download Now**. The **Download Status** dialog displays.
 - c. Click **Close** to close the Download status dialog.
 - d. Click **View Download Status**. A new browser window opens and displays the Download log. For example:


```
Download log:
Local database download at: 2008/08/11 17:40:42-0400
Downloading from ftp://1.1.1.1/list-1000000-cat.txt
Download size:      16274465
Database date: Sat, 09 Aug 2008 08:11:51 UTC
Total URL patterns: 1000000
Total categories:  10
```
8. Click **Apply**.

Note: Incremental updates are not available for the local database.

See Also

- ❑ ["Applying Policy"](#)
- ❑ ["Applying Policy to Categorized URLs"](#)
- ❑ ["More Policy Examples"](#)
- ❑ ["Defining Custom Categories in Policy"](#)

Section E: Configuring Internet Watch Foundation

The Internet Watch Foundation (IWF) is a non-profit organization that provides enterprises with a list of known child pornography URLs. The IWF database features a single category called **IWF-Restricted**, which is detectable and blockable using policy. IWF can be enabled along with other content filtering services. For information on IWF, visit their Web site at <http://www.iwf.org.uk>

For information on enabling the IWF database, see "Setting up a Web Content Filter" on page 392.

To download the IWF database, see "Downloading a Content Filter Database" on page 395.

See Also

- ❑ "Applying Policy"
- ❑ "Applying Policy to Categorized URLs"
- ❑ "More Policy Examples"
- ❑ "Defining Custom Categories in Policy"

Section F: Configuring a Third-Party Vendor

The third-party vendors supported on the ProxySG appliance are Internet Watch Foundation (IWF), Optenet, and Proventia. If you installed Surfcontrol, I-Filter, or Intersafe while on a previous version of SGOS, that configuration is maintained during an upgrade to SGOS 6.5.

Only Optenet and Proventia can be configured using the Management Console. Use the CLI to configure IWF, Surfcontrol, I-Filter, and Intersafe. SGOS 6.5 is the last branch to support these legacy providers.

The third-party vendor configuration tasks are identical and are covered in ["Setting up a Web Content Filter"](#).

See Also

- ❑ ["Applying Policy"](#)
- ❑ ["Applying Policy to Categorized URLs"](#)
- ❑ ["Defining Custom Categories in Policy"](#)

Section G: About Blue Coat Categories for YouTube

The appliance recognizes three types of YouTube URLs. The effectiveness of policy and coaching pages differs amongst the different URL types. Refer to the following table.

URL Type	Example(s) of User Action	Deny policy works	Coaching page works
YouTube homepage	User plays a video at www.youtube.com within a desktop web browser. The URL starts with: http://www.youtube.com/watch?v=	Yes	Yes
YouTube mobile website	User plays a video at www.youtube.com within a mobile web browser (the URL redirects to http://m.youtube.com/index). The URL starts with: http://m.youtube.com/watch?v=	Yes	Yes
Embedded in an <iframe>	User plays a video that has been embedded in a blog post. The URL could start with: http://www.youtube.com/embed/	Yes. The deny page is confined within the <iframe>.	Yes. The coaching page is confined within the <iframe>.
Video transport	User plays a YouTube playlist within a desktop web browser. The URL could start with: *.c.youtube.com/videoplayback?	Yes. The user may see an error message for the blocked video.	No

Note: In April 2015, Google deprecated YouTube Data API v2.0. As a result, Blue Coat categories for YouTube are no longer supported in SGOS versions prior to version 6.5.7.6.

Note: This feature is provided on an "as-is" basis. Symantec has no control of, and is not responsible for, information and content provided (or not) by YouTube. Customer is required to apply and use its own API key in order to activate this feature, and therefore obligated to comply with all terms of use regarding the foregoing (for example, see <https://developers.google.com/youtube/terms>),

including quotas, restrictions and limits on use that may be imposed by YouTube. Symantec shall not be liable for any change, discontinuance, availability or functionality of the features described herein.

For information on implementing coaching pages, refer to the “Notify User” action in *Visual Policy Manager Reference and Advanced Policy Tasks*.

The list of categories is static. In the Visual Policy Manager, you can view the categories in the category list (**Configuration > Edit Categories**) and in the Request URL Category object, but you cannot add, rename, edit, or remove them.

Setting the YouTube Server Key

Before you can enable YouTube as a provider, you must obtain an API server key and set it on the ProxySG appliance. For instructions, refer to the following article:

<http://www.symantec.com/docs/TECH241321>

After you set the server key, select YouTube in the Management Console at **Configuration > Content Filtering > General**.

Distinguishing Blue Coat Categories for YouTube in the Access Log

If the feature is enabled and categories are selected, and if you use an extended log file format (ELFF) for access logs, you can use the existing `category` access-log field to report on specified categories; however, the categories will be logged without the provider name. In addition, some categories share names with Blue Coat categories, such as Entertainment.

To distinguish between Blue Coat-defined categories and Blue Coat categories for YouTube in the access log, specify the ELFF field `cs-categories-qualified`. This field provides a list of all content categories of the request URL, qualified by the provider. For example, traffic matching YouTube’s Entertainment category would be logged as `Entertainment@YouTube`.

For information on access log formats, see [Chapter 30: "Access Log Formats"](#) on page 657.

Section H: Applying Policy

Even if you have enabled and downloaded a content filtering database on the ProxySG, you cannot regulate access to Web content until you create and install policy. This section discusses the interaction between content filtering categories and applications, and the creation and application of control policies.

Policy allows you to configure a set of rules that are used to filter Web content for HTTP, HTTPS, FTP, MMS and IM protocols. After you create and install policy on the ProxySG, every incoming client request is checked to determine if a rule matches the requested content. If a rule matches the request, the ProxySG uses the action specified in the rule to handle the incoming request.

Note: A URL can belong to a maximum of 16 categories. If a URL is assigned to more than 16 categories, the policy rules that you define will not apply for the additional categories.

Applying Policy to Categorized URLs

Policy rules are created to restrict, allow, and track Web access. Every content filter database provides pre-defined categories that you can reference in policy to create rules.

The examples in this section are created using the Visual Policy Manager (VPM) in the ProxySG. For composing policy using the Content Policy Language (CPL), refer to the *Blue Coat Content Policy Language Guide*.

The VPM layers that are relevant for configuring content filtering policy are:

- ❑ **Web Authentication Layer** (<Proxy> Layer in CPL)—Determines whether users must authenticate to the ProxySG for accessing Web content. If your content filtering policy is dependent on user identity or request characteristics, use this layer.
- ❑ **Web Content Layer** (<Cache> Layer in CPL)—Determines caching behavior, such as verification and ICAP redirection. If you are using content filtering to manage a type of content globally, create these rules in this layer.
- ❑ **Web Access Layer** (<Proxy> Layer in CPL)—Determines access privileges and restrictions for users when they access Web content.
- ❑ **SSL Access Layer** (<SSL> Layer in CPL)—Determines the allow or deny actions for HTTPS traffic.


Creating a Blacklist

If your default proxy policy is set to allow and you would like to block users access to certain categories, you must create policy to block all requests for the categories that you wish to restrict access in your network.


In this example, Sports/Recreation, Gambling, and Shopping categories are blocked with a single rule and a predefined exception page `content_filter_denied` is served to the user. This exception page informs the user that the request was denied because the requested content belongs to a category that is blocked.

To create a blacklist using VPM:


1. Select the **Configuration > Policy > Policy Options** tab.




2. Verify that the **Default Proxy Policy** option is set to **Allow**.
3. Access the VPM (**Configuration > Policy > Visual Policy Manager**).



4. Add a rule in a **Web Access Layer**:
 - a. In the **Destination** column, right click and select **Set**. The **Set Destination Object** dialog displays.
 - b. In the **Set Destination Object** dialog, click **New > Request URL Category**. The **Add Request URL Category Object** dialog displays.
 - c. Expand the list of categories for your content filter database in the **Categories** list.



5. Select the categories to block and click **OK**. This example blocks **Shopping**, **Gambling** and **Sports/Recreation** categories.



6. Set the action for blocking the categories In the **Action** column, right click and select **Deny** or **Deny Content Filter**.


The **Deny** action, denies the user access without providing an explanation for the denial of the requested content. And the **Deny Content Filter** action, denies the user access to the requested content and describes that the request was denied because it belongs to a category blocked by organizational policy.

Restricting Access by Category and Time of Day


When the default proxy policy is set to allow, the following example illustrates how to restrict access to sports during core business hours. In this example, access to sports is only allowed between noon and 1 pm, and then between 7 pm and midnight. This example uses a single rule and denies both blocks of time where access to this category is not permitted.

To create policy to restrict access by category and time:


1. Select the **Configuration > Policy > Policy Options** tab.




2. Verify that the **Default Proxy Policy** option is set to **Allow**.
3. Access the VPM (**Configuration > Policy > Visual Policy Manager**).




4. Add a rule in a **Web Access Layer**:
 - a. In the **Destination** column, right click and select **Set**. The **Set Destination Object** dialog displays.
 - b. In the **Set Destination Object** dialog, click **New > Request URL Category**. The **Add Request URL Category Object** dialog displays. In this example, the **Add Request URL Category Object** is named **Sports Access**.
 - c. Expand the list of categories for your content filter database from the **Categories** list.




5. Select the categories to block and click **OK**.




6. Set the time to block access to the category:
 - a. Select the **Time** column, right click and select **Set**. The **Set Time Object** dialog displays.
 - b. Select **New > Time** in the **Set Time Object** dialog. The **Add Time Object** dialog displays.



7. Configure the **Time Object**:
 - a. Select the **Time Zone** and **Enable** options. Add the time interval that you wish to restrict access to sports. This example restricts access between 1:05 pm and 7:00 pm.
 - b. Repeat step **a** to create a second time object to restrict access between midnight and noon.
8. Create a Time Combined Object:
 - a. Select **New > Combined Time Object** in the **Set Time Object**.
 - b. Create a **Combined Time Object** to add both time intervals (Time1 and Time2) in one rule.
9. Click **OK**.



a.




10. Set the action to restrict access. In the **Action** column, right click and select **Deny** or **Deny Content Filter**

The **Deny** action, denies the user access without providing an explanation for the denial of the requested content. And the **Deny Content Filter** action, denies the user access to the requested content and describes that the request was denied because it belongs to a category blocked by organizational policy.

Configuring Authentication-Based Access Privileges


Prerequisite: To configure access privileges using authentication, authentications realms must be configured on the ProxySG. Authentication realms allow you to create policy to exempt certain users or groups from accessing specified content while allowing access to specific individuals or groups.


The following example illustrates how to restrict software downloads to users in the IT group only. The default proxy policy in this example is **Allow**.



Note: While the following example blocks most downloads, it will not prevent all Web downloads. For example, compressed and encrypted files, server side scripts and Webmail attachments are not detected.


1. Add a rule in a **Web Authentication Layer** to authenticate users before granting access to Web content. This policy layer prompts the user for authentication:
 - a. In the **Action** column, right click and select **Set**. The **Set Action Object** dialog displays.
 - b. In the **Set Action Object** dialog, click **New > Authenticate**. The **Add Authenticate Object** dialog displays. Select the authentication mode and realm.
 - c. Click **OK** to save your changes and exit.
2. Add a rule in a **Web Access Layer** to restrict access to downloads by file extension and by Apparent Data Type of the content:
 - a. In the **Destination** column, right click and select **Set**. The **Set Destination Object** dialog displays.
 - b. In the **Set Destination Object** dialog, click **New > File Extensions**. The **Add File Extension Object** dialog displays.
 - c. In the **Known Extensions** field, find and add .exe files. Click **OK**.






- d. Select the apparent data types that include DOS and Windows executables and Windows Cabinet files. In the **Set Destination Object** dialog, click **New > Apparent Data Type**, and select the choices. Click **OK**.

Note: If your network environment includes either a **Content Analysis** or **ProxyAV** appliance, you can check **Enable ICAP Scanning** in the **New Apparent Data Type** object to have your ICAP server validate the data type of files contained in archive files (zip, rar).



- e. Combine the two rules using a combined object. In the **Set Destination Object** dialog, click **New > Combined Destination Object** and add the file extensions and the apparent data type rule created above. Click **OK**



3. Specify the desired action, by setting the action to **Deny**
4. Select the source field and click **New > Group**. Browse for the IT user group and click OK.
5. Right-click the source field in this rule and click **Negate**.


This rule will now prevent all users except for those in the Active Directory group, **IT** from being able to download the executables and CAB files.

Creating a Whitelist


If the default policy on the ProxySG is set to deny, you must create a whitelist to permit Web access to users. Whitelists require constant maintenance to be effective. Unless your enterprise Web access policy is very restrictive, Symantec recommends setting the default policy to allow. The default policy of allow will keep the help desk activity less hectic in managing Web access policies.

To create a whitelist using VPM:


1. Select the **Configuration > Policy > Policy Options** tab.




2. Verify that the **Default Proxy Policy** is **Deny**.



3. Add a rule in a **Web Access Layer**:
 - a. In the **Destination** column, right click and select **Set**. The **Set Destination Object** dialog displays.
 - b. In the **Set Destination Object** dialog, click **New > Request URL Category**. The **Add Request URL Category Object** dialog displays.
 - c. Expand the list of categories for your content filter database in the **Categories** list.



4. Select the categories to allow and click **OK**. This example allows **Business/Economy** and the **Computers/Internet** categories.




5. Set the action for blocking the categories In the **Action** column, right click and select **Allow**.


Creating Policy to Log Access to Specific Content

To monitor Web content requests from users in the network, you can record information in the ProxySG event log. For example, you can create policy to allow or deny access to a category and record users who attempt to access the specified category. The following example, illustrates how to use policy to track users who access the **Adult/Mature Content** category.

To log Web content access in an event log:



1. Add a rule in a **Web Access Layer**:
 - a. In the **Destination** column, right click and select **Set**. The **Set Destination Object** dialog displays.
 - b. In the **Set Destination Object** dialog, click **New > Request URL Category**. The **Add Request URL Category Object** dialog displays.
 - c. Expand the list of categories for your content filter database in the **Categories** list.
 - d. Select the categories to monitor and click **OK**. This example tracks access of **Adult/Mature Content**.



2. Select the information to be logged. This example logs information on the username, domain and IP address.
 - a. In the Web Access Layer, select the **Track** column and right-click and select **New > Event Log**.
 - b. Select from the list of **Substitution Variables** to log specific details about the **URL** or the **USER** and click **OK**. For information on substitution variables, refer to the *Visual Policy Manager Reference*.

Creating Policy When Category Information is Unavailable


An attempt to categorize a URL fails if no database is downloaded, your license is expired, or if a system error occurs. In such a case, the category is considered *unavailable* and triggers to block a category are not operative because the ProxySG is unable to determine the category. When the policy depends on the category of a URL, you do not want such errors to inadvertently allow ordinarily restricted content to be served by the ProxySG.

The category *unlicensed* is assigned in addition to *unavailable* when the failure to categorize occurred because of license expiry. This can be caused by the expiration of your Blue Coat license to use content filtering, or because of expiration of your license from the provider.


The following example illustrates how to block access (this is a mode of operation called *fail-closed*) to the requested content when category information is unavailable.

The **System** category *unavailable* includes the *unavailable* and *unlicensed* conditions. The *unlicensed* condition helps you identify that the category was not identified because the content filter license has expired.

To create policy when the category for a requested URL is unavailable:



1. Add a rule in a **Web Access Layer**:
 - a. In the **Destination** column, right click and select **Set**. The **Set Destination Object** dialog displays.
 - b. In the **Set Destination Object** dialog, click **New > Request URL Category**. The **Add Request URL Category Object** dialog displays.
 - c. Expand the **System** category list.
2. Select the category to monitor:
 - a. Select the checkbox for the **System** category **unavailable**.
 - b. Click **OK**.



3. Set the action to restrict access. In the **Action** column, right click and select **Deny Content Filter**.

You can also use this feature with custom exception pages (refer to the *Visual Policy Manager Reference*), where a custom exception page displays during business hours, say between 8 am and 6 pm local time for the requested content. In the event that the license is expiring, the user can be served an exception page that instructs the user to inform the administrator about license expiry.

Creating Policy for Uncategorized URLs

URLs that are not categorized are assigned the system category `none`. This is *not* an error condition; many sites (such as those inside a corporate intranet) are unlikely to be categorized by a commercial service. Use `category=none` to detect uncategorized sites and apply relevant policy. The following example disallows access to uncategorized sites outside of the corporate network:

```
define subnet intranet
  10.0.0.0/8 ; internal network
  192.168.123.45; external gateway
end
<proxy>
  ; allow unrestricted access to internal addresses
  ALLOW url.address=intranet

  ; otherwise (internet), restrict Sports, Shopping and
  uncategorized sites
  DENY category=(Sports, Shopping, none)
```

Creating Policy for Controlling Web 2.0 Applications

The Web 2.0 policy control objects are listed below:

- ❑ **Request URL Application:** The **Request URL Application** object gives you the ability to block popular Web applications such as Facebook, LinkedIn, or Pandora. As new applications emerge or existing applications evolve, WebFilter tracks the HTTP requests that these Web applications use to serve content, and provides periodic updates to include the new requestsdomains that are added. You can use the **Request URL Application** object to block an application and all the associated requests automatically.

For the applications you have blocked, you do not have to update your policy to continue blocking the new content sources; To block newly recognized applications, you will need to select the new applications and refresh your network policy.

- ❑ **Request URL Operation:** The **Request URL Operation** object restricts the actions a user can perform on a Web application. For instance, when you select the **Upload Picture** action for the **Request URL Operation**, you create a single rule that blocks the action of uploading pictures to any of the applications or services where the action can be performed such as Flickr, Picasa, or Smugmug.

When you block by operation, unlike blocking by application, you prevent users in your network from performing the specified operation for all applications that support that operation. They may however, be able to access the application itself.

Note, however, that the Request URL operation object only pertains to operations for sites that WebFilter recognizes as Web applications. So, blocking picture uploads would not prevent users in your network from using FTP to upload a JPEG file to an FTP server, or from using an HTTP POST to upload a picture on a Web site running bulletin board software.

Pre-requisites for Controlling Web Applications

- Proxy Edition license (not a MACH5 license)
- Symantec WebFilter license.
- The Symantec WebFilter feature must be enabled. (**Configuration > Content Filtering > General**)
- A current WebFilter database must be downloaded to the ProxySG. (**Configuration > Content Filtering > Blue Coat WebFilter**)
- The ProxySG must have one or more Web services, such as External HTTP and HTTPS, set to intercept. Bypassed Web traffic is not classified into applications.

Policy Examples Using the Application Control Objects

Use Case: Allow users to access Facebook and LinkedIn, but block access to other social networking sites. Also, block access to all games, including access to games on Facebook.

1. Launch the ProxySG Management Console.
2. Launch the Visual Policy Manager (VPM).

Select **Configuration > Policy > Visual Policy Manager**, and click **Launch**.

3. Create the rules to allow access to Facebook and LinkedIn, but restrict access to all other social networking sites. You must define the allow Facebook and LinkedIn rule before the rule that blocks access to other social networking sites.

To allow access to Facebook:

- a. Add a Web Access Layer. Select **Policy > Add Web Access Layer**.
- b. On the **Destination** column, right click and select **Request URL Application**.
- c. Select **Facebook** and **LinkedIn** from the application list and click **OK**.

Note: To filter through the list of supported applications, you can enter the name of the application in the **Filter applications by:** pick list. Based on your input, the on-screen display narrows the list of applications. You must then select the application(s) for which you want to create rules.

- d. Set **Action** to **Allow**.


To restrict access to all other social networking sites:

- a. Select **Edit > Add Rule** to add a new rule in the same Web Access layer.
- b. On the **Destination** column, right click and select **Request URL Category**.
- c. Select the **Social Networking** category from the list that displays and click **OK**.
- d. On the **Action** column, right click and select **Deny**. Your rules should look like this:



4. To properly block access to all games, including those on Facebook, you need to create another Web Access layer that defines the rule as follows:
 - a. Add a new Web Access Layer. Select **Policy > Add Web Access Layer**.

- b. On the **Destination** column, right click and select **Request URL Category**.
- c. Select the **Games** category from the list that displays and click **OK**.




5. Click **Install Policy**. You have now installed policy that blocks all games in your network, and permits access to the Facebook and LinkedIn applications in the social networking category.

Use Case: Allow limited access on Facebook but deny access all other sites in the social networking category. In this example, you restrict users from uploading attachments, videos or pictures on Facebook, but allow all other operations that the application supports.

1. Launch the ProxySG Management Console.
2. Launch the Visual Policy Manager (VPM).

Select **Configuration > Policy > Visual Policy Manager**, and click **Launch**.
3. Add a Web Access Layer. Select **Policy > Add Web Access Layer**.
4. Create a rule that allows access to Facebook but restricts uploads.
 - a. On the **Destination** column, right click and select **Set > New > Combined Destination Object**.
 - b. Select **New > Request URL Application** and select **Facebook** from the list of applications.
 - c. Select **New > Request URL Operation**.
 - d. Select **Upload Attachment, Upload Videos and Upload Pictures** from the list of operations and click **OK**.
 - e. Create the rule that checks for the application and the associated operation.
 - Select the application object you created for Facebook in Step 4b and **Add it to At least one of these objects**.
 - Select the **Negate** option in the bottom list. The display text changes from **AND At least one of these objects** to **AND None of these objects**. Then select the operation object you created for the uploading actions in Step 4d and click **Add**. Your policy should look as follows:



- Click **OK** to exit all open dialogs.
- f. On the **Action** column of the Web Access Layer, right click and select **Allow**.

You have now created a rule that matches on the application Facebook and prevents the action of uploading attachments, pictures or video. When a user attempts to upload these items on Facebook, the action will be blocked.
- 5. Restrict access to all other social networking sites.
 - a. Select **Edit > Add Rule** to add a new rule in the same Web Access layer.
 - b. On the **Destination** column, right click and select **Request URL Category**.
 - c. Select the **Social Networking** category from the list that displays and click **OK**.
 - d. On the **Action** column, right click and select **Deny**.
- 6. Click **Install Policy**. Your policy rule that allows limited access on Facebook and blocks access all other social networking sites is installed on the ProxySG appliance.

See Also

- ❑ *Content Policy Language Reference*
- ❑ *Command Line Interface Reference*

Verify the Application Filtering Policy is Working Properly

After you have installed your application filtering policy, you will want to verify that the policy works as you intended. From a client workstation on the network:

- ❑ Verify that you cannot access websites of blocked categories.
- ❑ Confirm that you can access websites of allowed categories.
- ❑ For each application that is unitarily blocked, verify that you cannot access any component of the application.
- ❑ For each operation that is blocked for an application, verify that you cannot perform that operation for that application. In addition, verify that you can perform operations that are not denied.

If your policy is not working properly, make sure you have spelled the application and operation names exactly as listed in the `view applications` and `view operations` command output. Also make sure that the operation is supported by the application. Correct any errors, install the revised policy, and run through the above verification steps again.

Additional Information

The following two access log variables are available in the Symantec Reporter access log format (bcreportermain_v1):

x-bluecoat-application-name

```
x-bluecoat-application-operation
```

Limitations

The policy compiler will not display a warning if you create policy that defines unsupported combinations of application names and operations. For example, Twitter doesn't support uploading of pictures but the compiler doesn't warn you that the following policy is invalid:

```
url.application.name=Twitter url.application.operation="Upload Pictures"
deny
```

More Policy Examples

Use Case: Limit employee access to travel Web sites.

The first step is to rephrase this policy as a set of rules. In this example, the model of a general rule and exceptions to that rule is used:

- ❑ Rule 1: All users are denied access to travel sites
- ❑ Rule 2: As an exception to the above, Human Resources users are allowed to visit Travel sites

Before you can write the policy, you must be able to identify users in the Human Resources group. You can do this with an external authentication server, or define the group locally on the ProxySG. For information on identifying and authenticating users, see ["Controlling User Access with Identity-based Access Controls"](#) on page 898 and for information on authentication modes supported on the ProxySG see ["About Authentication Modes"](#) on page 908.

In this example, a group called `human_resources` is identified and authenticated through an external server called `my_auth_server`.

This then translates into a fairly straightforward policy written in the local policy file:

```
<proxy>
; Ensure all access is authenticated
Authenticate(my_auth_server)
<proxy>
; Rule 1: All users denied access to travel
DENY category=travel
<proxy>
; Rule 2: Exception for HR
ALLOW category=travel group=human_resources
DENY category=sites
```

Use Case: Student access to Health sites is limited to a specified time of day, when the Health 100 class is held.

This time the policy contains no exceptions:

- ❑ Rule 1: Health sites can be accessed Monday, Wednesday, and Friday from 10-11am.
- ❑ Rule 2: Health sites can not be accessed at other times.

```
define condition Health_class time
    weekday=(1, 3, 5)  time=1000..1100
end
<proxy>
; 1) Allow access to health while class in session
    ALLOW category=health condition=health_class_time
; 2) at all other times, deny access to health
    DENY category=health
```

Defining Custom Categories in Policy

Custom categories give administrators the ability to create their own filtering criteria. This ability allows administrators to create specific categories that lists Web sites and keywords to block or allow and can be adapted to their organizational requirements.

Custom categories are created in the policy file using the VPM or CPL. If you have extensive category definitions, Blue Coat recommends that you put them into a local database rather than into a policy file. The local database stores custom categories in a more scalable and efficient manner, and separates the administration of categories from policy. See "[Configuring a Local Database](#)" on page 405.

To add URLs to a category, you only need to specify a partial URL:

Note: The local database produces only the most specific URL match and returns a single category.

The same policy syntax will produce a different match. If more than one category is provided, policy processing may match more than one category and hence will return more than one category. See "[Local Database Matching Example](#)" on page 406 for more information.

- ❑ hosts and subdomains within the domain you specify will automatically be included
- ❑ if you specify a path, all paths with that prefix are included (if you specify no path, the whole site is included) For example, if you add
www2.nature.nps.gov/air/webcams/parks/grcacam/nps.gov/grca
only the pages in the /grca directory of nps.gov will be included in the category, but if you just add www2.nature.nps.gov/ all pages in the entire directory will be included in the category.

Note: If a requested HTTPS host is categorized in a content filtering database, filtering rules apply even when HTTPS Intercept is disabled on the ProxySG. However, if the request contains a path or query string and the categorization relies on the host/relative path, the categorization results could be different. This is because the path or query is not accessible when HTTPS Intercept is disabled. The difference in categorization is caused as a result of categorizing the host name only versus using the host name and path or query string.

Example:

```
define category Grand_Canyon
  kaibab.org
  www2.nature.nps.gov/air/webcams/parks/grcacam
  nps.gov/grca
  grandcanyon.org
end
```

Any URL at `kaibab.org` is now put into the `Grand_Canyon` category (in addition to any category it might be assigned by a provider). Only those pages in the `/grca` directory of `nps.gov` are put in this category.

Nested Definitions and Subcategories

You can define subcategories and nest category definitions by adding a `category=<name>` rule. To continue the example, you could add:

```
define category Yellowstone
  yellowstone-natl-park.com
  nps.gov/yell/
end
define category National_Parks
  category=Grand_Canyon; Grand_Canyon is a subcategory of
National_Parks
  category=Yellowstone; Yellowstone is a subcategory of National_Parks
  nps.gov/yose; Yosemite - doesn't have its own category (yet)
end
```

With these definitions, pages at `kaibab.org` are assigned *two* categories: `Grand_Canyon` and `National_Parks`. You can add URLs to the `Grand_Canyon` category and they are automatically added by implication to the `National_Parks` category as well.

Multiple unrelated categories can also be assigned by CPL. For example, by adding:

```
define category Webcams
  www2.nature.nps.gov/air/webcams/parks/grcacam
end
```

the URL, `http://www2.nature.nps.gov/air/webcams/parks/grcacam/grcacam.htm`, will have three categories assigned to it:

- ❑ `Grand_Canyon` (because it appears in the definition directly)
- ❑ `National_Parks` (because `Grand_Canyon` is included as a subcategory)
- ❑ `Webcams` (because it also appears in this definition)

However, the other sites in the `Grand_Canyon` category are not categorized as `Webcams`. This can be seen by testing the URL (or any other you want to try) clicking the **Test** button on the Management Console.

You can test for any of these categories independently. For example, the following example is a policy that depends on the above definitions, and assumes that your provider has a category called `Travel` into which most national park sites probably fall. The policy is intended to prevent access to travel sites during the day, with the exception of those designated `National_Parks` sites. But the `Grand_Canyon` webcam is an exception to that exception.

Example:

```
<proxy>
  category=Webcams DENY
  category=National_Parks ALLOW
  category=Travel time =0800..1800 DENY
```

Click the **Test** button on the Management Console or the `test-url` command in CLI to validate the categories assigned to any URL. This can help you to ensure that your policy rules have the expected effect (refer to *Configuring Policy Tracing* in the *<Emphasis>Blue Coat Content Policy Language Guide*).

If you are using policy-defined categories and a content-filter provider at the same time, be sure that your custom category names do not coincide with the ones supplied by your provider. You can also use the same names—this adds your URLs to the existing categories, and extends those categories with your own definitions. For example, if the webcam mentioned above was not actually categorized as Travel by your provider, you could do the following to add it to the Travel category (for the purpose of policy):

```
define category Travel ; extending a vendor category
  www2.nature.nps.gov/air/webcams/parks/grcacam/ ; add the GC webcam
end
```

Note: The policy definitions described in this section can also be used as definitions in a local database. See "[Configuring a Local Database](#)" on page 405 for information about local databases.

Section I: Troubleshooting

This section describes troubleshooting tips and solutions for content filtering issues. It discusses the following topics:

- ❑ "Unable to Communicate with the WebPulse Service" on page 440
- ❑ "Event Log Message: Invalid WebPulse Service Name, Health Check Failed" on page 440
- ❑ "Error Determining Category for Requested URL" on page 441
- ❑ "Error Downloading a Content Filtering Database" on page 442

Unable to Communicate with the WebPulse Service

Symantec WebFilter and WebPulse are enabled, and the following error message displays:

```
Dynamic categorization error: unable to communicate with service 0  
510000:1 ../protocols/cerberian/Cerberian_api.cpp:79
```

To resolve this issue:

1. Use DNS to resolve `sp.cwfservice.net`.

Note: The ProxySG appliance resolves the domain name `sp.cwfservice.net` once a day and maintains the list of returned IP addresses. The ProxySG appliance then uses the IP address that provides the fastest service. If an IP address that is in use fails to respond, the ProxySG appliance will failover to an alternate IP address. Health checks are automatically conducted on all the IP addresses to make this failover as smooth as possible and to restore service to the geographically closest IP address as soon as it is available.

2. Check the firewall logs for messages about denied or blocked traffic attempting to reach IP addresses or in response from IP addresses. A firewall rule denying or blocking in either direction impedes WebPulse.

Event Log Message: Invalid WebPulse Service Name, Health Check Failed

The following event log message displays:

```
Invalid WebPulse service name - Health check failed - Receive failed.
```

These messages are common in event logs and, for the most part, should not affect your service. A server may fail an L4 health check for various reasons, but unless all servers (services) are unavailable for extended periods of time, you should not experience interruptions in WebPulse services and can regard this as expected behavior.

When the proxy makes a request for the WebPulse service name, several IP addresses for our servers are returned. The ProxySG appliance will periodically perform a quick Layer-4 health check (opening and closing a TCP socket with no data transfer) to each of those servers. In the event that the ProxySG appliance cannot contact the server or doesn't receive a response quickly enough, it logs similar event log messages.

Your WebPulse service will not be interrupted unless all of the servers are unable to be contacted for more than a few seconds. When one of these error messages appears, the services health status changes back to healthy within 2 to 10 seconds.

Error Determining Category for Requested URL

The access log shows the category for a URL as **Unavailable**; Category **Unavailable** indicates that an error occurred when determining the category for a requested URL.

The following is an example access log message:

```
2007-08-07 22:19:02 59 10.78.1.98 404 TCP_NC_MISS 412 428 GET http
www.sahnienterprise.com 80 /images/menu.gif - - - DIRECT
www.sahnienterprise.com text/html;%20charset=iso-8859-1 http://
www.sahnienterprise.com/Mozilla/5.0 (Windows; U; Windows NT 5.1; en-
US; rv:1.8.1.6) Gecko/20070725 Firefox/2.0.0.6 PROXIED "Unavailable" -
10.78.1.100
```

Start by manually testing the URL in the **URL** field in **Configuration > Content Filtering > General** on the Management Console.

If the category is still **Unavailable**, go through the list of possible causes in the following table.

<i>Possible Causes</i>	<i>Check the Following</i>
The database is not installed	Check <code>show content-filter status</code> .
The database is corrupt.	Check <code>show content-filter status</code> .
The database has expired.	Check the validity of the database. To verify that the latest content filter database is available on the ProxySG, enter the following commands in the CLI: Blue Coat SG210 Series> en Enable Password: XXXXX Blue Coat SG210 Series# show content-filter status Provider: Blue Coat Status: Ready
A communication error occurred contacting the WebPulse service.	Check the event log entries for WebPulse messages.
The ProxySG appliance license has expired.	If you are using a trial or demo license, instead of a perpetual license, the ProxySG license may have expired. Verify the status of your license on the Maintenance > Licensing > View tab. To purchase a license, contact Symantec Technical Support or your Symantec sales representative.
(Possible, but not likely) There are issues with memory or a disk error.	Check event log entries for disk or memory messages.

Error Downloading a Content Filtering Database

To view the status of your database download, click **View Download Status** on the **Configuration > Content Filtering > Vendor_Name** tab.

- ❑ For the **ERROR: HTTP 401 - Unauthorized**, verify that you have entered your username and password correctly. For example, the following error message was generated when an incorrect username was entered to download a Optenet database:

```
Download log:
Optenet download at: Thu, 21 June 2007 18:03:08
Checking incremental update
Checking download parameters
Fetching:http://example.com/
Warning: HTTP 401 - Unauthorized
Downloading full control file
Optenet download at: Thu, 21 June 2007 18:03:17
Downloading from http://example.com/
Fetching:http://example.com/
ERROR: HTTP 401 - Unauthorized
Download failed
Download failed
Previous download:
...
```

If you have an upstream proxy and all internet traffic must be forwarded to this upstream proxy, you must enable `download-via-forwarding` on this ProxySG using the following CLI command:

```
SGOS> enable
SGOS# config t
SGOS#(config) forwarding
SGOS#(config forwarding) download-via-forwarding enable
```

- ❑ For the **Socket Connection Error**, check for network connectivity and Internet access in your network.

Only after completing network troubleshooting, perform the following procedure if the socket connection error persists.

Because the content filter database is downloaded using SSL, if the SSL client on the ProxySG gets corrupt, a connection error occurs.

1. Verify that you have a valid SSL client on the ProxySG.
 - a. Access the Command Line Interface (CLI) of the ProxySG appliance.
 - b. In configuration mode, view the SSL client configuration.

```
Blue Coat SG210 Series>en
Enable Password:XXXXX

Blue Coat SG210 Series#conf t
Blue Coat SG210 Series#(config) ssl
Blue Coat SG210 Series#(config ssl)view ssl-client
SSL-Client Name      Keyring      CCL          Protocol
default              <None>      browser-trusted TLSv1.2vTLSv1.1
```


2. If you have an ssl-client configured but the issue still persists, delete, and recreate the SSL client.

- a. In the Configuration mode:

```
Blue Coat SG210 Series#(config ssl)delete ssl-client  
ok
```

```
Blue Coat SG210 Series#(config ssl)create ssl-client default  
defaulting protocol to TLSv1.2vTLSv1.1 and CCL to browser-trusted  
ok
```


Chapter 21: Configuring Threat Protection

Blue Coat ProxySG and the Symantec threat-protection appliances work in conjunction to analyze incoming Web content and protect users from malware and malicious content. Malware is defined as software that infiltrates or damages a computer system without the owner's informed consent. The common types of malware include adware, spyware, viruses, downloaders and Trojan horses.

Symantec's threat protection solution protects user productivity, blocks malware downloads and Web threats, and enables compliance to network security policies.

The following sections describe how to configure threat protection with the internal Content Analysis service or the ProxySG and Content Analysis / ProxyAV appliances:

- ❑ ["About Threat Protection"](#)
- ❑ ["Enabling Malware Scanning"](#)
- ❑ ["Updating the Malware Scanning Policy"](#)
- ❑ ["Fine Tuning the Malware Scanning Policy using VPM"](#)
- ❑ ["Disable Malware Scanning"](#)
- ❑ ["Edit an ICAP Content Analysis Service"](#)
- ❑ ["Delete an ICAP service From the List of ICAP services"](#)

About Threat Protection

Owing to the interactive nature of the Internet, enterprises are constantly exposed to Web threats that can cause damage to company data and productivity. To ensure that your users, systems and data are protected at all times, Symantec provides a multi-layered solution in a single appliance that protects you from existing and emerging threats.

Content Analysis is an advanced module that, when used in conjunction with the SGOS Proxy module, encompasses all facets of network-level threat protection:

- ❑ Web Filtering protection during policy execution using Symantec Web Filtering services,
- ❑ Anti-virus scanning with multiple vendors
- ❑ File Whitelisting to reduce resource load on known-good files
- ❑ Sandboxing with Symantec MAA and FireEye external appliance solutions, to analyze suspicious files, and update WebPulse with the results.

Symantec's threat protection solution is a cohesive solution that provides the intelligence and control required to manage Web traffic in your network. It includes Content Analysis or ProxyAV that provides in-path threat protection from malware, and the Symantec WebFilter and WebPulse service that provide URL filtering and a Web-based community watch service.

WebPulse is a 'community watch' cloud service that detects hidden malware and provides reputation and Web content analysis in real time. WebPulse services are offered to all customers who have a valid Symantec WebFilter license. For more information on WebPulse, see "About Symantec WebFilter and the WebPulse Service" on page 381.

In addition to providing reputation and Web categorization information, the WebPulse service proactively notifies all Symantec WebFilter subscribers of emerging malware threats. This notification is possible because of the malware feedback mechanism between the ProxySG and Symantec's ICAP analysis services, Content Analysis and ProxyAV.

The ProxySG monitors the results of content scans and notifies the WebPulse service when a new virus or malware is found. This notification triggers an update of the Symantec WebFilter database and all members of the WebPulse community are protected from the emerging threat.

Symantec's threat protection solution also provides a threat protection policy that is implemented when you integrate the appliances and enable malware scanning. The malware scanning policy that is implemented is predefined set of policies that offer optimal protection. The Malware Scanning policy can be set to optimize either your network security needs or your network performance needs.

Internal ProxyAVThreat Protection Configuration Tasks

An ICAP service to leverage the internal Content Analysis service is configured automatically on the Advanced Secure Gateway appliance. Unlike with external ICAP services detailed in the proceeding table, this internal ICAP service does not require any additional configuration or licensing steps.

External ProxyAVThreat Protection Configuration Tasks

The tasks that must be completed for configuring threat protection between the ProxySG and the ProxyAV or Content Analysis service are listed in the following table.

Table 21–1 Tasks for Configuring Threat Protection

Task	Task Description
1. Install and configure the ICAP appliance.	- Configure the ProxyAV or Content Analysis service with basic network settings. Make sure to configure the ICAP server and the ProxySG appliances on the same subnet. From the ProxyAV Management Console, perform the following tasks: - Activate the ICAP server licenses, as appropriate. - Configure the scanning behavior on the ProxyAV or Content Analysis Management Console For information on these tasks, refer to the <i>ProxyAV Configuration and Management Guide</i> or the <i>Content Analysis System WebGuide</i>.
2. Select whether to transfer data between the ProxySG and the ProxyAV or Content Analysis service using plain ICAP or secure ICAP.	The ProxySG and the ProxyAV or Content Analysis appliances can communicate with each other using plain ICAP, secure ICAP or both methods. If you wish to use secure communication mode between appliances, either use the built-in SSL device profile or create a new SSL device profile to authorize ProxyAV or Content Analysis on the ProxySG. For information on SSL device profile, see "About SSL Device Profiles" on page 1291. If you create an SSL device profile, make sure that the CA certificate is imported in the ProxySG at Configuration > SSL > External Certificates. Otherwise, when the Verify Peer option is enabled in Configuration > SSL > Device Profiles, the ProxySG fails to verify ProxyAV or Content Analysis as trusted. For information on enabling secure connection on ProxyAV or Content Analysis, or creating a new certificate, refer to the <i>ProxyAV Configuration and Management Guide</i> or the <i>Content Analysis System WebGuide</i>.
3. Add the ProxyAV to allow in-path threat detection and enable malware scanning, on the ProxySG.	To add the ProxyAV to the ProxySG, see "Adding an ICAP service for Content Scanning". To begin scanning of Web responses you must enable malware scanning. Malware scanning, when enabled, automatically invokes a predefined threat protection policy. See "Enabling Malware Scanning".


Adding an ICAP service for Content Scanning

Symantec ICAP services, (ProxyAV and Content Analysis) are designed to prevent malicious content from entering your network. When you add these services to your configuration, the ProxySG redirects Web responses fetched from the origin Web server to the ICAP service to be scanned before delivering the content to the user.

The protocol that the ProxyAV or Content Analysis and the ProxySG use to communicate is the Internet Content Adaptation Protocol or ICAP.

To for content scanning:

1. Select **Configuration> Threat Protection> Malware Scanning**
- 2.
3. Select **New**. The Add New CAS/ProxyAV ICAP Server dialog displays.



4. In the **Server host name or IP address** field, enter the host name or IP address of the ICAP server. Only an IPv4 address is accepted.
5. Choose the connection mode(s) and ports. The default is plain ICAP only. If you select secure ICAP, you must add an SSL device profile. An SSL device profile contains the information required for device authentication, including the name of the keyring with the private key and certificate this requires to be authenticated. For information on SSL device profiles, see ["About SSL Device Profiles"](#).
6. Click **OK** to save your changes and exit the open dialog box. You now have **proxyavx** service that is automatically created to perform response modification. Response modification means that the ICAP service only acts on requested content that is redirected to it by the ProxySG after the content is served by the origin Web server.
7. Click **Perform health check** to verify that the ICAP server is accessible. The health check result is displayed immediately. For information on health checks, see ["Managing ICAP Health Checks"](#).
8. Continue with ["Enabling Malware Scanning"](#).

Enabling Malware Scanning

To begin content scanning after adding an ICAP service to the ProxySG, Symantec provides a built-in threat protection policy with a set of predefined rules. These rules protect your network from malicious content while supporting your network performance or network protection needs.

Enabling malware scanning implements the threat protection policy on the ProxySG. The threat protection policy compiles policy conditions based on your preferences in the malware scanning configuration. By default, when you enable malware scanning, the options selected in the malware scanning configuration supports high performance scanning using a secure ICAP connection between the ICAP service and the ProxySG, if available, and the user is denied access to the requested content if the scan cannot be completed for any reason.

The rules used by the predefined threat protection policy can be made more or less strict with the use of a simple radio button configuration. That is, while the threat-protection policy itself does not change, only conditions that match your configuration settings are implemented from the threat protection policy file. And, when you change configuration, the compiled policy is automatically updated to reflect the configuration changes.

Note: The threat protection policy cannot be edited. If you would like to supplement or override the configuration in this policy, see ["Fine Tuning the Malware Scanning Policy using VPM"](#).

To enable malware scanning:

1. Go to **Configuration > Threat Protection > Malware Scanning**. By default, malware scanning is disabled on the ProxySG.
2. Verify that one or more ICAP services are added for content scanning. For information on adding a ICAP service, see ["Adding an ICAP service for Content Scanning"](#).
3. Select the **Enable malware scanning** checkbox. The threat protection policy is invoked with the malware scanning options selected in configuration or the pre-set default.
4. Click **Apply** to save your changes.
5. (Optional) To modify the malware scanning options that constitute the rules in the threat protection policy for your network, see the following topics:
 - ["Selecting the Protection Level"](#)
 - ["Selecting the Connection Security Mode"](#)
 - ["Defining the Scan Failure Action"](#)

6. (Optional) To verify that malware feedback to the WebPulse service is enabled, check that the **Enable WebPulse service** checkbox is selected in **Configuration > Threat Protection > WebPulse**. By default, when Symantec WebFilter is enabled on the ProxySG, WebPulse is enabled.
For information on WebPulse, see "[About Symantec WebFilter and the WebPulse Service](#)". For information on configuring WebPulse, see "[Configuring WebPulse Services](#)".

Selecting the Protection Level


The threat protection policy offers two levels for scanning ICAP responses — high performance and maximum security. While the ProxyAV scans all Web responses when set to maximum security, it selectively scans Web responses when set to high performance bypassing content that has a low risk of malware infection.

The high performance option is designed to ensure network safety while maintaining quick response times for enterprise users. For example, file types that are deemed to be low risk, such as certain image types, are not scanned when set to high performance. To view the content that is not scanned with the high performance option, in configuration mode of the CLI enter `show sources policy threat-protection`.

The scanning rules configured for high performance and maximum security are subject to change, as Symantec may update rules based on the latest Web vulnerabilities and security risk assessments. To obtain the latest version of the malware scanning policy, see "[Updating the Malware Scanning Policy](#)" on page 453.

To set the protection level:

1. Select **Configuration > Threat Protection > Malware Scanning**.
2. Select the **Protection level** preference for your enterprise. The default protection level is **High performance**.
3. Click **Apply** to save your changes.



For information on adding rules in VPM to make an exception to the configured protection level, see ["Fine Tuning the Malware Scanning Policy using VPM"](#).


Selecting the Connection Security Mode

The communication between the ProxySG and external ICAP servers can be in plain ICAP, secure ICAP or can use both plain and secure ICAP, depending on whether the response processed by the ProxySG uses the HTTP, FTP, or HTTPS protocol.

Plain ICAP should be used only for non-confidential data. In particular, if plain ICAP is used for intercepted HTTPS traffic, then data intended to be cryptographically secured would be transmitted in plain text on the local network. With secure ICAP data exchange occurs through a secure data channel. This method protects the integrity of messages that are sent between the ProxySG and the external ICAP server.

To select a connection security mode:

1. Select **Configuration > Threat Protection > Malware Scanning**.



2. Select the **Connection security** preference for your network:
 - a. **Always use secure connections** ensures that all communication between the ProxySG and the ICAP server uses SSL-encrypted ICAP. By default, secure ICAP uses port 11344.
 - b. **Use secure connections for encrypted requests, if available** is the default option and it ensures that requests will be sent over secure ICAP, if the service supports it.
 - c. **Always use plain connections** sets all communication between the ProxySG and the external ICAP service in non-secure mode. This option is available only if the service object is configured to support plain ICAP.
 - d. Click **Apply** to save your changes.

Defining the Scan Failure Action

If an error occurs while scanning a file, you must configure the ProxySG for the action that it must take. The action you define allows the ProxySG to either serve the requested content to the client or deny the request when the scan cannot be completed.

A scan might fail because the ICAP service is not available due to a health check failure, a scanning timeout, a connection failure between the ProxySG and the ICAP server, or an internal error on the ICAP server. For maximum security, the recommended and default setting is to deny the request when failures occur.

In addition to setting the action on an unsuccessful scan globally, you can configure policy for individual ICAP scanning errors. For information on ICAP error scan codes, see ["Editing an ICAP Service"](#) on page 481.

The rule is configured only to determine the action in the event an error occurs while scanning a file.

To select an action upon an unsuccessful scan:

1. Select **Configuration > Threat Protection > Malware Scanning**

2. Select your preferred failure option under **Action on an unsuccessful scan**. To ensure network security, the default is **Deny the client request**.
3. Click **Apply** to save your changes.

Viewing the Installed Malware Scanning Policy

After configuring the options for malware scanning, you can view the policy that was compiled and installed on the ProxySG using the Management Console.

To view the installed policy using the Management Console:

1. Select **Configuration > Policy > Policy Files**
2. Select **Current Policy** in the **View Policy** drop-down menu
3. Click **View**. The current policy installed on your ProxySG displays in a new window.

The malware scanning policy begins with the following text:

```
<Cache BC_malware_scanning_solution>
  policy.BC_malware_scanning_solution
...
```

See Also

- ❑ ["Updating the Malware Scanning Policy"](#)
- ❑ ["Fine Tuning the Malware Scanning Policy using VPM"](#)


Updating the Malware Scanning Policy

Because the threat landscape changes rapidly, Symantec facilitates updates to the threat protection solution to protect your network from the latest malware attacks and exploits. The threat protection policy updates are independent of SGOS upgrades.

Updates to the threat protection solution are available as a gzipped tar archive file which can be downloaded to a local Web server in your network or installed directly on the ProxySG.

To update the malware scanning policy directly on the ProxySG:

1. Select **Configuration > Threat Protection > Malware Scanning**.



2. Click **Update malware scanning policy**. The **Install Malware Scanning Policy** dialog box displays.
3. (Optional) Enter the **Installation URL**. Otherwise, accept the default URL.
If you have downloaded the threat protection policy to a local Web server, add the URL for the local Web server in this field.

Note: If you change the default URL, you cannot revert to the default value. You must manually re-enter the URL.

4. Click **Install**.
5. (Optional) Click **View** to view the contents of the updated threat protection policy file.

Note: The threat protection policy cannot be edited.

6. Click **OK** to save your changes and exit.

Fine Tuning the Malware Scanning Policy using VPM

When malware scanning is enabled, the threat protection policy file is invoked. The rules implemented in the threat protection policy either use the defaults or the selections that you configured in the malware scanning options in **Configuration > Threat Protection > Malware Scanning**.

Unlike other policy files, the threat protection policy file is not displayed in the **Policy Evaluation Order** list in **Policy > Policy Options > Policy Options** and the threat protection policy file cannot be edited or modified. However, you can create rules in the local policy file or in VPM policy to supplement or override the configured defaults. The rules created in local or VPM policy supersede the configuration in the threat protection policy because of the evaluation order of policy files. By default on the ProxySG, policy files are evaluated in the following order — Threat protection, VPM, Local, Central, and Forward.


The threat protection policy is evaluated first to provide you with the flexibility to adapt this policy to meet your business needs. For example, even if the malware scanning mode is configured at maximum protection through configuration, you can create rules in VPM to allow all traffic from internal hosts/subnets to be scanned using the high performance mode. Alternatively, if the default malware

scanning mode is high performance, you can add rules in VPM to invoke maximum protection mode for sites that belong to select content filtering categories such as software downloads or spyware sources.


The following example demonstrates how to create rules in VPM to complement the malware scanning options that are set in configuration. The setting in configuration, in the example below, uses maximum security. The VPM rule allows internal traffic to be scanned using the high performance rules that are defined in the threat protection policy.

Example: Configure high performance scanning for internal traffic


1. Set the Scanning mode in **Configuration > Threat protection > Malware Scanning** to **Maximum protection**.
2. Launch the VPM and create policy to scan all traffic from an internal host using the high performance mode. This example uses the 10.0.0.0/8 subnet.
 - a. Select **Configuration > Policy > Visual Policy Manager**.
 - b. Click **Launch**. The Visual Policy Manager displays in a new window.
 - c. Select **Policy > Add Web Content Layer**.
 - d. In the **Action** column, right click and select **Set**. The Set Action Object dialog box displays.
The first option in the Set Action Object dialog, **Always Verify**, is highlighted by default, but it is not an active selection. Continue with the next step.
 - e. In the **Set Action Object** dialog, click **New > Set Malware Scanning**. The Add Malware Scanning Object dialog displays.
 - f. Select **Perform high performance malware scan**.



- g. Click **OK** to save your changes and exit all open dialogs.



- h. In the **Destination** column, right click and select **Set**. The Set Destination Object dialog box displays.
 - i. Select **Destination IP Address/Subnet**. The Set Destination IP/Subnet Object dialog displays.
 - j. Add the IP address and subnet for the internal host and click **Close**.
 - k. Click **OK** to save your changes and exit all open dialogs.
 - l. Click **Apply** to install the policy. After this policy is installed, all traffic from the internal subnet 10.0.0.0/8 will be scanned using the high performance mode.
3. The completed rule is shown below.



Disable Malware Scanning

If you prefer to manually create policy for content scanning rather than use Symantec's threat protection solution that provides pre-defined rules for ICAP-based malware scanning, follow the instructions below.

Note: Malware scanning cannot be disabled if the threat protection solution is referenced in policy. For example, if you have created a rule in the Web Content layer that references the threat protection policy file, disabling malware scanning will cause policy compilation to fail. You must remove all references to the threat protection policy file before disabling malware scanning.

To disable malware scanning:

1. Select **Configuration > Threat Protection > Malware Scanning**.
2. Clear the **Enable malware scanning** checkbox.
3. Click **Apply**.
4. For information on creating policy, refer to the *Visual Policy Manager Reference*.

Edit an ICAP Content Analysis Service

A Content Analysis service is a collection of attributes that defines the communication between the ProxySG and external or internal ICAP services such as Content Analysis, ProxyAV or DLP services.

Use the following procedure to edit the service URL, change the maximum number of connections, modify ICAP service ports, or set ICAP options.

To edit a ProxyAV service:

1. Select **Configuration > External Services > ICAP > ICAP Services**.
2. Select the service to edit.
3. Click **Edit**. The **Edit ICAP Service** dialog displays.
4. Edit the service options, as desired (service URL, maximum number of connections, ICAP service ports, or ICAP options).
5. Click **OK**.
6. Click **Apply** to save your changes.
7. (Optional) For modifying advanced configuration options, see "[Editing an ICAP Service](#)" on page 481.


Delete an ICAP service From the List of ICAP services

Use the following steps to remove an ICAP service from the list of configured Content Analysis servers.

Note: If you have enabled malware scanning and have added only one Content Analysis or ProxyAV appliance for content scanning, you must disable malware scanning before you can delete that service from the **CAS/ProxyAV ICAP Servers** list. Malware scanning must be disabled so that the ICAP service to be deleted is no longer referenced in the threat protection policy.

To delete an external CAS or ProxyAV ICAP service:

1. Select **Configuration > Threat Protection > Malware Scanning**.
2. Select the ICAP service to be deleted from the **CAS/ProxyAV ICAP Servers** list.



3. Click **Delete**; click **OK** to confirm.
4. Click **Apply**. The service is deleted from the **CAS/ProxyAV ICAP Servers** list.

Chapter 22: Malicious Content Scanning Services

This chapter describes how to configure the ProxySG to interact with Internet Content Adaptation Protocol (ICAP) servers to provide content scanning.

The ProxySG appliance supports ICAP connections with external Content Analysis, ProxyAV, Symantec DLP, and other third party ICAP services.

To integrate external Content Analysis or ProxyAV with the ProxySG, see ["Configuring Threat Protection"](#).

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ [Section A: "About Content Scanning"](#) on page 460
- ❑ [Section B: "Configuring ICAP Services"](#) on page 473
- ❑ [Section C: "Securing Access to an ICAP Server"](#) on page 488
- ❑ [Section D: "Monitoring ICAP Requests and Sessions"](#) on page 496
- ❑ [Section E: "Creating ICAP Policy"](#) on page 503
- ❑ [Section F: "Managing Virus Scanning"](#) on page 517

Section A: About Content Scanning

Internet Content Adaptation Protocol (ICAP) is an open standard protocol that allows content engines to send HTTP based content to an ICAP server for performing value added services.

An ICAP server can filter, modify, or adapt Web content to meet the needs of your enterprise. The Advanced Secure Gateway, when integrated with a supported ICAP server such as Content Analysis and ProxyAV, provides content scanning, filtering, and repair service for Internet-based malicious code, in addition to reducing bandwidth usage and latency.

To eliminate threats to the network, the ProxySG forwards a Web request and/or the response to the ICAP server. The ICAP server filters and adapts the requested content, based on your needs, then returns the content to the ProxySG. The scanned and adapted content is then served to the user who requested the content, and stored on the ProxySG object store. For frequently accessed Web content, this integrated solution provides defense against malware, along with the benefits of limiting bandwidth usage and latency in the network.

Plain ICAP and Secure ICAP

The transaction between the ProxySG and the ICAP server can be executed using plain ICAP, secure ICAP or both. Plain ICAP is useful for scanning non-confidential data (HTTP).

Secure ICAP is SSL encrypted ICAP and requires an SSL license; both the ProxySG and the ICAP server must support secure ICAP. While Secure ICAP can be used for both HTTP and HTTPS traffic, plain ICAP is faster than secure ICAP because it does not have to deal with any encryption overhead. Therefore, Symantec recommends that you only use secure ICAP when scanning confidential data.

Content Processing Modes

An ICAP server processes Web content that is directed to it during Proxy policy evaluation. The content that the ICAP server receives can be processed in two modes — request modification and response modification.

Request modification (REQMOD)—Allows modification of outbound client requests. These requests are sent from the ProxySG to the ICAP server on their way to the origin content server. This is represented in the Visual Policy Manager as Perform Request Analysis.

Response modification (RESPMOD)—Allows modification of inbound client requests. These requests are sent from the ProxySG to the ICAP server after the requested content is retrieved from the origin content server.

REQMOD or RESPMOD is an attribute that is specified in the ICAP service, which is configured between the ProxySG and the ICAP server. This is represented in the Visual Policy Manager as Perform Response Analysis.

About Response Modification

The ProxySG sends the first part (a preview) of the object to the ICAP server that supports response modification. The object preview includes the HTTP request and response headers, and the first few bytes of the object. After checking those bytes, the ICAP server either continues with the transaction (that is, asks the ProxySG to send the remainder of the object for scanning) or sends a notification to the appliance that the object is clean and opts out of the transaction.

The response modification mode enables scanning of HTTP responses, remote system file retrieval or FTP RETR responses, FTP over HTTP, and SSL-intercepted response data.

Returning the Object to the ProxySG Appliance

For response modification, the returned object can be the original unchanged object, a repaired version of the original object minus a virus, or an error message indicating that the object contained a virus. Each of these responses is configured on the ICAP server, independent of the appliance and the ICAP protocol. If the appliance receives the error message, it forwards the error message to the client and does not save the infected file.

The following diagram illustrates the response modification process flow.




Figure 22-1 Response Modification Process Flow

About Request Modification

Request modification means the ICAP server scans contents that a client is attempting to send outside the network. This prevents unaware users from forwarding corrupted files or Webmail attachments. Request modification is also a method of content filtering and request transformation, which is used to protect network identification. Based on the results of the scan, the server might return an HTTP response to the client (for example, sports not allowed); or the client request might be modified, such as stripping a referer header, before continuing to the origin content server.

Request modification mode enables scanning of HTTP GET requests, PUT requests and POST requests, FTP upload requests and outgoing Webmail.

Note: Some ICAP servers do not support virus scanning for request modification, but support only content filtering.

The following diagram illustrates the request modification process flow.




Figure 22–2 Request Modification Process Flow

Caching and Serving the Object

After an object has been scanned and is determined to be cacheable, the ProxySG caches it and serves it for a subsequent request. When the appliance detects that the cached content has changed on the origin server, it fetches a fresh version, then forwards it to the ICAP server for scanning. If the ProxySG uses policies in the ICAP configuration, the policy applies to content fetches, distributions, refreshes, and pipelined requests.

For more information on policies, see [Section E: "Creating ICAP Policy"](#). For more information on the <Cache> layer, refer to the <Emphasis>Blue Coat Content Policy Language Guide.

ICAP v1.0 Features

This section describes the options for the ICAP v1.0 protocol that are provided on the ProxySG.

Sense Settings

The Sense Settings feature allows the ProxySG to query any identified ICAP server running v1.0, detect the parameters, and configure the ICAP service as appropriate. See "[Creating an ICAP Service](#)" on page 474.

ISTags

ISTags eliminates the need to designate artificial pattern version numbers, as was required in v0.95.

Every response from an ICAP v1.0 server must contain an IStag value that indicates the current state of the ICAP server. For instance, when the pattern/scanner version of a virus scanner on the ICAP server changes, the IStag value changes. This change invalidates all content cached with the previous IStag value and a subsequent request for any content in cache must be refetched from the origin content server and scanned by the ICAP server.

Backing out a virus pattern on the ICAP server can revert ISTags to previous values that are ignored by the ProxySG. To force the ProxySG to recognize the old values, use the Sense Settings option. See "[Creating an ICAP Service](#)" on page 474.

Persistent Connections

New ICAP connections are created dynamically as ICAP requests are received (up to the defined maximum connection limit). The connection remains open to receive subsequent requests. If a connection error occurs, the connection closes to prevent more errors.

Determining Which Files to Scan

In determining which files to scan, this integrated solution uses the content scanning server's filtering in addition to Proxy capabilities. The following table describes the supported content types and protocols.

Table 22–1 Content Types Scanned By ICAP Server and the ProxySG

ICAP Server supported content types	Proxy supported protocols	Unsupported content protocols
All or specified file types, based on the file extension, as configured on the server. Examples: .exe (executable programs), .bat (batch files), .doc and .rtf (document files), and .zip (archive files); or specific MIME types.	• All HTTP objects (uploaded or downloaded) • All FTP over HTTP (webftp) objects (uploaded or downloaded) • All native FTP objects (uploaded or downloaded) The above is true for both transparent and explicit proxies.	• Streaming content (for example, RTSP and MMS) • Live HTTP streams (for example, HTTP radio streams) • CIFS • MAPI • IM • TCP tunnel traffic
	HTTPS connections terminated at a ProxySG	HTTPS connections tunneled through a ProxySG

Whenever an object is requested or being refreshed and it was previously scanned, the Proxy verifies whether the pattern file has been updated since it was last scanned. If it was, the object is scanned again, even if the content has not changed. If the content has changed, the object is rescanned.

With the Proxy, you can define flexible, yet enterprise-specific content scanning policies, which are discussed in the following two sections.

Improving the User Experience

Object scanning adds another operation to the user process of requesting and receiving Web content. Therefore, the user might experience extremely slight noticeable delays during Web browsing as ICAP servers scan content. The ProxySG allows you to mitigate slower browse times and educate your users about what is occurring on their systems. This section discusses:

- ❑ Patience pages
- ❑ Data trickling
- ❑ Deferred scanning and infinite streams

About Patience Pages

Patience pages are HTML pages displayed to the user if an ICAP content scan exceeds the specified duration (seconds). You can configure the content of these pages to include a custom message and a help link. Patience pages refresh every five seconds and disappear when object scanning is complete.

Notes

- ❑ Patience pages are not compatible with *infinite stream* connections—or live content streamed over HTTP—such as a cam or video feed. ICAP scanning cannot begin until the object download completes. Because this never occurs with this type of content, the ProxySG continues downloading until the maximum ICAP file size limit is breached. At that point, the ProxySG either returns an error or attempts to serve the content to the client (depending on fail open/closed policy). However, even when configured to fail open and serve the content, the delay added to downloading this large amount of data is often enough to cause the user give up before reaching that point.
- ❑ Patience pages are limited to Web browsers.

About Data Trickling

Patience pages provide a solution to appease users during relatively short delays in object scans. However, scanning relatively large objects, scanning objects over a smaller bandwidth pipe, or high loads on servers might disrupt the user experience because connection time-outs occur. To prevent such time-outs, you can allow *data trickling* to occur. Depending on the trickling mode you enable, the ProxySG either trickles—or allows at a very slow rate—bytes to the client at the beginning of the scan or near the very end.

The ProxySG begins serving server content *without* waiting for the ICAP scan result. However, to maintain security, the full object is not delivered until the results of the content scan are complete (and the object is determined to not be infected).

Note: This feature is supported for the HTTP proxy only; FTP connections are not supported.

Trickling Data From the Start

In *trickle from start* mode, the ProxySG buffers a small amount of the *beginning* of the response body. As the ICAP server continues to scan the response, the ProxySG allows one byte per second to the client.




Figure 22-3 A client receives only the initial bytes of a transaction during the ICAP scan.

After the ICAP server completes its scan:

- ❑ If the object is deemed to be clean (no response modification is required), the ProxySG sends the rest of the object bytes to the client at the best speed allowed by the connection.
- ❑ If the object is deemed to be malicious, the ProxySG terminates the connection and the remainder of the response object bytes—which in this case are the majority of the bytes—are not sent to the client.

Deployment Notes

- ❑ This method is the more secure option because the client receives only a small amount of data pending the outcome of the virus scan.
- ❑ One drawback is that users might become impatient, especially if they notice the browser display of bytes received. They might assume the connection is poor or the server is busy, close the client, and restart a connection.

Trickling Data at the End

In *trickle at end* mode, the ProxySG sends the response to the client at the best speed allowed by the connection, except for the last 16 KB of data. As the ICAP server performs the content scan, the ProxySG allows one byte per second to the client.




Figure 22-4 A client receives most of the bytes immediately during the ICAP scan.

After the ICAP server completes its scan, the behavior is the same as described in ["Trickling Data From the Start"](#) on page 467.

Deployment Notes

- ❑ Symantec recommends this method for media content, such as flash objects.
- ❑ This method is more user-friendly than trickle at start. This is because users tend to be more patient when they notice that 99% of the object is downloaded versus 1%, and are less likely to perform a connection restart. However, network administrators might perceive this method as the less secure method, as a majority of the object is delivered before the results of the ICAP scan.

Deciding between Data Trickling and Patience Pages

ProxySG configuration options plus policy allow you to provide different ICAP *feedback* actions depending upon the type of traffic detected:

- ❑ Blue Coat defines interactive as the request involving a Web browser. Web browsers support data trickling and patience pages.
- ❑ Non-interactive traffic originates from non-browser applications, such as automatic software download or update clients. Such clients are not compatible with patience pages; therefore, data trickling or no feedback are the only supported options.

Based on whether the requirements of your enterprise places a higher value either on security or availability, the ProxySG allows you to specify the appropriate policy. However, you must also consider the user agents involved when determining the appropriate feedback method. For example, streaming clients cannot deliver patience pages, but they are susceptible to connection time-outs. Therefore, trickling is the suggested method. The following diagram provides basic guidelines for deciding which feedback method to implement.




Figure 22–5 Deciding which ICAP feedback method to employ.

Recommendations for Proxy Chaining Deployments

Proxy chaining deployments are common in enterprises, especially in core/branch office scenarios. Data trickling is achievable, but behavior is dependent upon how the ProxySG appliances are configured. The following are common deployment scenarios.

- ❑ **The downstream ProxySG is performing ICAP scanning, and the upstream ProxySG is not:** Data trickling and patience pages are not affected in this scenario.
- ❑ **The upstream ProxySG is performing ICAP scanning, and the downstream ProxySG is not:** The only issue with this deployment is that user agent-specific policy cannot be applied at the core ProxySG because the branch ProxySG consolidates multiple client requests in one out-going request to the upstream ProxySG. If data trickling is employed at the upstream ProxySG and if ICAP scanning detects a virus, the upstream ProxySG resets the client connection. This also deletes the corrupted object from the downstream ProxySG cache.
- ❑ **Both ProxySG appliances (upstream and downstream) are scanning:** Behavior is mostly determined by the configuration of the upstream ProxySG.
 - If the upstream ProxySG is configured to deliver patience pages, then the downstream ProxySG also attempts to serve patience pages, including to non-graphical user agents. Therefore, this method is not recommended.
 - If the upstream ProxySG employs data trickle from start, the downstream ProxySG is not able to send any bytes to the client for a long period of time. If a patience page is not configured on the downstream ProxySG, users might experience connection time-outs.
 - If the upstream ProxySG employs trickle at end, the downstream ProxySG allows for all options of patience page and data trickling.

Avoiding Network Outages due to Infinite Streaming Issues

Infinite streams are connections such as web cams or flash media—traffic over an HTTP connection—that conceivably have no end. Characteristics of infinite streams may include no content length, slow data rate and long response time. Because the object cannot be fully downloaded, the ICAP content scan cannot start; however, the connection between the ProxySG and the ProxyAV appliance remains, which wastes finite connection resources.

The deferred scanning feature (enabled by default) solves the infinite streaming issue by detecting ICAP requests that are unnecessarily holding up ICAP connections (without requiring the ProxyAV appliance) and defers those requests until the full object has been received.

How Deferred Scanning Works

Deferred scanning detects the possibility of infinite streams by the fact that the number of ICAP resources in use has reached a certain threshold. It then defers the scanning of those streams by deferring the oldest, outstanding ICAP requests first. For every new ICAP request, the ProxySG does the following:

- ❑ If the total number of outstanding ICAP actions for the current server has reached the defer threshold, the ProxySG defers the oldest ICAP connection that has not yet received a full object.

The defer threshold is specified by the administrator as a percentage. For example, if the defer threshold is set to 70 percent and the maximum connections are set to 100, then up to 70 connections are allowed before the ProxySG begins to defer connection which have not finished downloading a complete object.

Note: See ["Creating an ICAP Service"](#) on page 474 for information about setting the defer scanning threshold value on the ProxySG Management Console.

When an ICAP connection is deferred, the connection to the ICAP server is closed. The application response continues to be received and when the download is complete the ICAP request is restarted. The new ICAP request may still be queued if there are no available ICAP connections. After a request is deferred, ICAP waits to receive the full object before restarting the request. If there is a queue when a deferred action has received a complete object, that action is queued behind other deferred actions that have finished. However it will be queued before other new requests.

Deferred Scanning and Setting the Feedback Options

Depending on how you configure the ICAP feedback option (patience page or data trickling) and the size of the object, deferred scanning might cause a delay in ICAP response because the entire response must be sent to the ICAP server at once. The feedback option allows you to specify the type of feedback you want to receive during an ICAP scan. For information about setting feedback options, see ["Configuring ICAP Feedback"](#) on page 482.

If a patience page is configured, the browser continues to receive a patience page until the object is fully received and the outstanding ICAP actions have completed.

If the data trickle options are configured, the object continues to trickle during deferred scanning. However, because of the trickle buffer requirement, there might be a delay, with or without deferred scanning, before the ProxySG starts sending a response.

About ICAP Server Failover

When creating an ICAP action, you can specify a list of ICAP servers or groups to use, in order of preference. If the first server or group in the list does not pass the health checks, the ProxySG moves down the list until it finds a server or group that is healthy and uses that to perform the scanning.

The primary server resumes ICAP processing when the next health check is successful; the standby server or server group does not retain the primary responsibility.

Notes

- ❑ Failover is configured as part of the ICAP policy definition.
- ❑ You cannot configure failover policy until ICAP services are configured on the ProxySG.
- ❑ To avoid errors, ICAP service names cannot be named **fail_open** or **fail_closed** (the CLI commands prevent these names from being created).
- ❑

Section B: Configuring ICAP Services

This section describes how to configure the Proxy to communicate with an ICAP server for content scanning.

To configure threat protection with an external ProxyAV or Content Analysis, see ["Configuring Threat Protection"](#).

Overview of Configuring ICAP on the ProxySG

Table 3-2 provides a high-level view of workflow tasks for configuring Proxy/ICAP communications. It also provides task descriptions.

Table 22–2 Workflow Tasks—Configuring ProxySG ICAP Communications

Task	Task Description
1. Install and configure the ICAP server	Follow the manufacturer instructions for installing the ICAP server, including any configuration necessary to work with the ProxySG. Based on your network environment, you might use the ProxySG with multiple ICAP servers or multiple scanning services on the same server. Configure options as needed, including the exception message displayed to end users in the event the requested object was modified or blocked.
2. Decide whether to scan data using plain ICAP or secure ICAP	Scan data using the plain ICAP method, secure ICAP method or both. • Plain ICAP should be used only for non-confidential data. In particular, if plain ICAP is used for intercepted HTTPS traffic, then data intended to be cryptographically secured would be transmitted in plain text on the local network. • Secure ICAP send data through a secure data channel. This method protects the integrity of messages that are sent between the ProxySG and the ICAP server while it allows users to authenticate ICAP servers by enabling certificate verification.

Table 22–2 Workflow Tasks—Configuring ProxySG ICAP Communications (Continued)

Task	Task Description
3. (Optional—secure ICAP only) Select the default SSL profile or create an SSL device profile on the ProxySG appliance	An SSL device profile is required to authorize the ICAP server, if you use secure ICAP. For information on SSL device profile, see "About SSL Device Profiles". Note: If you create an SSL device profile, instead of using a built-in device profile, ensure that the ICAP server certificate is installed as trusted under External certificates. Otherwise, when the Verify Peer option is enabled, the ProxySG appliance fails to verify the ICAP server as a trusted server.
4. Create and configure new or existing ICAP services on the ProxySG. For information on ICAP content processing modes, see "Content Processing Modes" .	Create an ICAP service that specifies the ICAP server IP address, supported connection method, content processing mode and select deferred scanning, if desired. See "Creating an ICAP Service" .
5. Specify the feedback method	Select patience pages or data trickling for feedback method. See "Configuring ICAP Feedback" .
6. Add ICAP rules to policy	Depending on your network needs, add ICAP rules to policy and install the policy file on the ProxySG. See Section E: "Creating ICAP Policy"

Creating an ICAP Service

An ICAP service is a collection of attributes that defines the communication between the ProxySG and the ICAP server. It includes the server IP address or hostname, ICAP scanning method, and a host of other options including the supported number of connections.

Note: The internal Content Analysis service automatically creates both a request and response modification service. These services cannot be configured, but are available for use in the Set Request or Set Response Analysis service dialog in the Visual Policy Manager, as well as in the Malware Scanning configuration.


You must create an ICAP service for each ICAP server or scanning service. For example, if you are using the ProxySG with multiple ICAP servers or multiple scanning services (RESPMOD or REQMOD) on the same server, add an ICAP service for each server and RESPMOD or REQMOD service.

Similar ICAP scanning services can then be grouped together to create a service group that helps distribute and balance the load of scanning requests. Further, each ICAP service or service group can be accessed through VPM or CPL to configure policy for better administrative control.


The following instructions describe how to create an ICAP service for any supported third-party ICAP server.

To create and configure an ICAP service:

1. Select **Configuration > External Services> ICAP > ICAP Services**.



2. Add a new service:
 - a. Click **New**; the Add List Item dialog displays.
 - b. Enter an alphanumeric name in the **Add ICAP Service** field. This example uses **Request1**.
 - c. Click **OK**. The new ICAP object displays in the services list.
3. Highlight the ICAP service name and click **Edit**. The Management Console displays the Edit ICAP Service dialog.



4. Configure the service communication options:

Note: The default ICAP version is 1.0 and cannot be changed.

- a. In the **Service URL** field, enter the ICAP server URL, which includes the URL schema, ICAP server hostname or IP address. For example:
`icap://10.x.x.x/`. Check with your ICAP vendor for the appropriate URL format.
- b. (Introduced in 6.5.9.14) Identify the ICAP service type as being Threat Protection, DLP, or Other. Use this service if you are migrating appliance policy to the Symantec Web Security Service; however, policy rules that reference an Other service are not enforceable in the Web Security Service.
- c. In the **Maximum Number of Connections** field, enter the maximum possible connections at any given time that can occur between the ProxySG and the ICAP server. The range is a number from 5 to 4096. The default is 5. The number of recommended connections depends on the capabilities of the ICAP server. Refer to the vendor's product information.
- d. In the **Connection timeout** field, enter the number of seconds the ProxySG waits for replies from the ICAP server. This timeout is the duration for which the TCP connection between the ProxySG and the ICAP server is maintained. It helps verify the responsiveness of the ICAP server and prevents users from experiencing unnecessary delays. The default timeout is 70 seconds, and you can enter a value in the range of 1 to 65535.

Note: The connection timeout value does not measure how much of the scanning process is complete, it is a mechanism for ensuring that the communication between the appliances is alive and healthy. The details of the interaction between the ProxySG appliance and the ICAP server can only be viewed through a packet capture.

If the ICAP server does not respond within the configured timeout value, by default, the user will not receive the requested content. However, if Content Analysis or ProxyAV is your ICAP server, the scanning response configured in the **Configuration > Threat Protection > Malware Scanning** determines whether or not the user is served the requested content.

- e. Select **Defer scanning at threshold** to set the threshold at which the ProxySG defers the oldest ICAP connection that has not yet received a full object. The range is 0 percent – 100 percent. By default, the deferred scanning threshold is enabled when an ICAP service is created. The defer threshold scanning defaults to 80 percent.

- f. Select **Notify administrator when virus detected** to send an e-mail to the administrator if the ICAP scan detects a virus. The notification is also sent to the Event Log and the Event Log e-mail list.
- g. Select **Use vendor's "virus found" page** to display the default vendor error exception page to the client instead of the ProxySG exception page.

This is the default behavior for SGOS upgrades from previous versions. This feature maintains the same appearance of previous versions, but also retains the inherent timestamp issues involved with cache hits. If this option is not selected, the exception pages originate from the ProxySG, and they employ the accurate timestamps for cache hits.

5. Configure service ports for plain ICAP and secure ICAP. You can enable one or both types of ICAP connections at the same time. However, you must select at least one type of ICAP service.
 - a. Select **This service supports plain ICAP connections** to use plain ICAP. Use plain ICAP when you are scanning plain data (HTTP). In this case, if the HTTPS proxy is enabled on the ProxySG, the data is decrypted first on the ProxySG and then sent to the ICAP server.
 - b. In the **Plain ICAP port** field, enter a port number. The default port is 1344.
 - c. Select **This service supports secure ICAP connections** to use secure ICAP. Use secure ICAP when you are scanning sensitive or confidential data (HTTPS).
 - d. In the **Secure ICAP port** field, enter a port number. The default port is 11344.
 - e. If you selected secure ICAP, make sure that you select a valid SSL profile for secure ICAP in the **SSL Device Profile** field. This associates an SSL device profile with the secure ICAP service.

Note: If you do not select an SSL device profile you cannot use secure ICAP connections. The SSL device profile can be customized for your environment. For more information, see ["Appliance Certificates and SSL Device Profiles"](#) on page 1290.

6. Configure ICAP v1.0 features:

- a. Click **Sense Settings** to automatically configure the ICAP service using the ICAP server parameters.
- b. Select the ICAP method: response modification or request modification. This selection cannot be modified for an ICAP service created using the ["Adding an ICAP service for Content Scanning"](#) on page 447.

Note: An ICAP server might have separate URLs for response modification and request modification services.

- c. (Only for RESPMOD service) If you are using file scanning policies based on file extensions on the ProxyAV appliance, enter **0** in the **Preview size (bytes)** field, and select **enabled**. With a **0** bytes preview size, only response headers are sent to the ICAP server; more object data is only sent if requested by the ICAP server.

or

If you have enabled the Kaspersky Apparent Data Types feature on the ProxyAV appliance, enter a value (**512** is recommended) in the **Preview size (bytes)** field, and select **enabled**. The ICAP server reads the object up to the specified byte total. The ICAP server either continues with the transaction (that is, receives the remainder of the object for scanning) or opts out of the transaction.

or

Unselect **enabled** if the above two situations don't apply to you; do not use the preview option.

- d. (Optional) The **Send** options allow additional information to be forwarded to the ICAP server. Select one or more of the following: **Client address**, **Server address**, **Authenticated user**, or **Authenticated groups**.

- e. Click **Perform health check** to perform an immediate health check on this service.
 - f. Click **OK** to close the dialog.
7. Click **Apply**.

See Also


- ❑ ["About Content Scanning" on page 460](#)
- ❑ ["Configuring ICAP Services" on page 473](#)
- ❑ ["Avoiding Network Outages due to Infinite Streaming Issues" on page 470](#)
- ❑ ["Configuring ICAP Services" on page 473](#)
- ❑ ["Securing Access to an ICAP Server" on page 488](#)
- ❑ ["Monitoring ICAP Requests and Sessions" on page 496](#)
- ❑ ["Managing Virus Scanning" on page 517](#)

Managing ICAP Health Checks

ProxySG health check features allow you to perform tasks such as immediate checking, disable health checks, and override various notifications and settings.

To manage ICAP health checks:

1. Select **Configuration > Health Checks > General**.



2. Select an ICAP service or service group.
3. Click **Perform health check** to get an immediate connection status for the ProxyAV appliance or service group.
4. Click **Edit** to display the **Edit ICAP Health Check** dialog.

5. Select the **Enabled state**:
 - **Enabled**: Marks the ICAP service or group as enabled and functioning.
 - **Disabled: Healthy**: Marks the ICAP service as healthy, but not able to receive connections. One reason to select this option is to preserve current statistics; the disabled state is temporary.
 - **Disabled: Unhealthy**: Marks the ICAP service as down and not able to receive connections. One reason to select this is that you are taking the server offline for maintenance or replacement.
6. For a service group, select **All**, **Any** or *a number* from the drop-down menu to indicate the **Minimum number of members that must be healthy** for the service group to be considered healthy.
7. Click **Apply**.

For detailed information about the health check configuration options, including override features, see ["Configuring Global Defaults"](#) on page 1363.

Configure Alert Notifications for ICAP

You can set up alert notifications for queued and deferred ICAP connections.

To configure alert notifications for ICAP:

1. Select **Maintenance > Health Monitoring > General**.
2. Click the **General** tab.
3. Set alert notification properties for queued ICAP connections:
 - a. Select **ICAP Queued Connections** and click **Edit**.
 - b. In the dialog that appears, enter values for the Critical Threshold and Critical Interval. (Warning Threshold and Interval are not measured for this metric.)

See ["Planning Considerations for Using Health Monitoring"](#) on page 1338 for more information.
 - c. Select a notification output (**Log**; **Trap**; **Email**).
 - d. Click **OK**.
4. Set alert notification properties for deferred ICAP connections:
 - a. Select **ICAP Deferred Connections** and click **Edit**.
 - b. Repeat steps [3b](#) through [3d](#).
5. Click **Apply**.

Monitoring ICAP Health Metrics

When the ProxySG appliance powers on, both of the ICAP connections metrics are below threshold. If a threshold is exceeded for the duration specified by the interval value, the metric changes from **OK** to **Critical** and the new status is logged.

In order for a metric to be healthy again, it must return to and stay below threshold for the duration of the interval. When this happens, the new status is logged.

Example

The following example depicts the changing health of an ICAP connection metric configured with the default threshold (80%) and interval (120 seconds) and when log entries are created during the health monitoring process:

- ❑ Metric health starts at **OK**.
- ❑ The metric exceeds 80% for 60 seconds and then returns below threshold. The state is still **OK** and no log entry is created.
- ❑ The metric exceeds 80% again. After being above threshold for 120 seconds, the metric becomes **Critical** and a log entry is created for the Critical state.
- ❑ The metric goes below 80% for 100 seconds before exceeding the threshold. The state is still **Critical** and no log entry is created.
- ❑ The metric goes below threshold. After being below threshold for 120 seconds, the metric becomes **OK** and a log entry is created for the OK state.

Deleting an ICAP Service

The following steps describe how to delete an ICAP service.

Note: You cannot delete an ICAP service used in an ProxySG policy (that is, if a policy rule uses the ICAP service name) or that belongs to a service group. Before proceeding with the steps below, make sure to remove the references in policy and remove the ICAP service from the service group.

To delete an ICAP service to a third-party ICAP server:

1. Select **Configuration > External Services > ICAP > ICAP Services**.
2. Select the service to be deleted.
3. Click **Delete**; click **OK** to confirm.
4. Click **Apply**.

Editing an ICAP Service

The instructions below are for modifying the settings for the ICAP service configured between the ProxySG and Content Analysis, the ProxyAV appliance, or any third party ICAP server.

To edit the ICAP service:

1. Go to **Configuration > External Services > ICAP > ICAP Services**.
2. Continue with step 3 of "[Creating an ICAP Service](#)" on page 474.

Configuring ICAP Feedback

This section describes how to specify what type of feedback is provided to users during an ICAP scan. See ["Improving the User Experience"](#) on page 465.

To specify and configure the ICAP feedback method:

1. Select **Configuration > External Services > ICAP > ICAP Feedback**.

The screenshot shows the 'ICAP Feedback' configuration page. It has three tabs: 'ICAP Services', 'ICAP Feedback', and 'ICAP Patience Page'. The 'ICAP Feedback' tab is selected. The page is divided into two main sections: 'ICAP Feedback for Interactive Traffic' and 'ICAP Feedback for Non-Interactive Traffic'. Each section contains three radio button options. In the 'Interactive Traffic' section, the first option is 'Do not provide feedback during ICAP scanning' (labeled 2a), the second is 'Provide feedback after 5 seconds' (labeled 2b), and the third is a group of three options: 'Return patience page' (labeled 2c), 'Trickle object data from start', and 'Trickle object data at end'. The 'Non-Interactive Traffic' section has similar options, with the third option being 'Trickle object data at end' (labeled 3c).

2. Configure options for interactive traffic (browser-based requests):
 - a. The **Do not provide feedback...** option means that if users experience delays in receiving content, they are not notified as to the reason (ICAP scanning). Selecting this option greys out the other options.
 - b. The default duration to wait before notifying a client that an ICAP scan is occurring is five seconds. You can change this value in the **Provide feedback after** field, but if you make the value too long, users might become impatient and manually close the client, believing the connection is hung.
 - c. Select the feedback method:
 - **Return patience pages:** The client displays a Web page to the user providing a description of the delay (ICAP scanning). This page is customizable, as described in the next section.

Note: When the deferred scanning option is enabled and a patience page is configured, the browser continues to receive a patience page until the object is fully received and the outstanding ICAP actions have completed.

-
- **Trickle object data from start:** The client receives 1 byte per second, which should prevent connection time-outs while the ICAP server performs the scan. If the response from the ICAP server is clean, the client receives the rest of the object data at the best connection speed possible. If the scan detects malicious content, the connection is dropped. This is the more secure method.
 - **Trickle object data at end:** The client receives most (99%) of the object data, but the final bytes are sent at the rate of one per second while the ICAP scanner performs the scan. If the response from the ICAP server is clean, the client receives the rest of the object data at the best connection speed possible. If the scan detects malicious content, the connection is dropped. This is the least secure method, as most of the data has already been delivered to the client. However, this method provides the best user experience because most of the object is already delivered.

Note: When deferred scanning is enabled and the data trickle options are configured, the object continues to trickle during deferred scanning. However, due to the trickle buffer requirement, there may be a delay before the ProxySG starts sending a response.

3. Configure options for non-interactive traffic (content such as flash animation over HTTP):
 - a. The **Do not provide feedback...** option means that if users experience delays in receiving content, they are not notified as to the reason (ICAP scanning). Selecting this option greys out the other options.
 - b. The default duration to wait before notifying a client that an ICAP scan is occurring is five seconds. You can change this value in the **Provide feedback after** field, but if you make the value too long, users might become impatient and manually close the client, believing the connection is hung.
 - c. Select the feedback method:
 - **Trickle object data from start:** See the descriptions in Step 2.
 - **Trickle object data at end:** See the descriptions in Step 2.

4. Click **Apply**.

These configurations are global. You can define further feedback policy that applies to specific user and conditional subsets. In the VPM, the object is located in the Web Access Layer: **Return ICAP Feedback**.

Customizing ICAP Patience Text

This section describes how to customize text displayed during ICAP scanning. Patience pages are displayed if the appropriate option is selected, as described in the previous section: ["Improving the User Experience"](#) on page 465.

The following topics describe how to customize the HTTP/FTP patience page:


- ❑ "HTTP Patience Text" on page 484
- ❑ "FTP Patience Text" on page 487

HTTP Patience Text


The ProxySG allows you to customize the patience page components and text that are displayed to users when HTTP clients experience delays as Web content is scanned.

To customize HTTP patience pages:


1. Select **Configuration > External Services> ICAP > ICAP Patience Page**.



2. In the **HTTP Patience Page Customization** section, click **Header**, **Summary**, **Details**, or **Help**. The corresponding customize dialog displays. Customize the information as appropriate.




- a. **Custom Patience Header**—Contains HTML tags that define what displays in the dialog title bar. This component also contains the `<meta http-equiv>` tag, which is used to specify a non-English character set.



- b. **Custom Patience Summary Message**—HTML and text that informs users that a content scan is occurring.



- c. **Custom Patience Details Message**—Uses data to indicate scanning progress. The information includes the URL currently being scanned, the number of bytes processed, and the elapsed time of the scan.



- d. **Custom Patience Help Message**—Displays instructions for users should they experience a problem with the patience page.

3. Click **Apply**.

All of these components are displayed on the patience page.

Windows XP, Service Pack 2 Behavior

Microsoft is continually updating Windows XP security measures, which impacts how the ProxySG manages patience pages.

- ❑ Browsers running on Windows XP, Service Pack 2 (XP SP2), experience slightly different patience page behavior when pop-up blocking is enabled.
 - If pop-up blocking is not enabled, patience page behavior should be normal.
 - If pop-up blocking is enabled (the default), the ProxySG attempts to display the patience page in the root window.
 - If the download triggers an invisible Javascript window, the user can track the scanning progress with the progress bar at the bottom of the window; however, if other policy blocks Javascript active content, this bar is also not visible.
- ❑ If Internet Explorer blocks all downloads initiated by Javascript, the user must click the yellow alert bar to download the scanned object.
- ❑ Users experience two patience page responses for non-cacheable objects.

Interactivity Notes

- ❑ When ICAP scanning is enabled and a patience page is triggered, a unique URL is dynamically generated and sent to the browser to access the patience page. This unique URL might contain a modified version of the original URL. This is expected behavior.
- ❑ Patience pages and exceptions can only be triggered by left-clicking a link. If a user right-clicks a link and attempts to save it, it is not possible to display patience pages. If this action causes a problem, the user might see browser-specific errors (for example, an Internet *site not found* error); however, ICAP policy is still in effect.
- ❑ A patience page is not displayed if a client object request results in an HTTP 302 response and the ProxySG pipelines the object in the `Location` header. After the ProxySG receives the client request for the object, the client enters a waiting state because a server-side retrieval of the object is already in progress. The wait status of the client request prevents the patience page from displaying. To prevent the ProxySG from pipelining these requests (which decreases performance) and to retain the ability to provide a patience page, configure HTTP as follows:

```
#SGOS (config) http no pipeline client redirects
```
- ❑ The status bar update does not work if it is disabled or if the Javascript does not have sufficient rights to update it.


- ❑ **Looping:** Certain conditions cause browsers to re-spawn patience pages. For example, a site states it will begin a download in 10 seconds, initiates a pop-up download window, and returns to the root window. If the download window allows pop-ups, the patience page displays in a separate window. The automatic return to the root window initiates the download sequence again, spawning another patience page. If unnoticed, this loop could cause a system hang. The same behavior occurs if the user clicks the back button to return to the root window. For known and used download sites, you can create policy that redirects the page so that it doesn't return to the root window after a download starts.

FTP Patience Text

For content over FTP, the patience text displayed to FTP clients during an ICAP scan can be modified.

To customize FTP patience text:

1. Select **Configuration > External Services > ICAP > ICAP Patience Page**.



2. In the **FTP Patience Page Customization** field, click **Summary**; the Customize FTP Patience Text dialog displays. Customize the FTP client patience text as appropriate.
3. Click **OK**.
4. Click **Apply**.

Section C: Securing Access to an ICAP Server

You can secure access between the ProxySG appliance and an ICAP server using a variety of methods. Choosing the appropriate method is contingent upon your platform considerations and network topology.

Secure ICAP can be used regardless of your network topology or platform considerations. Because secure ICAP has no such restrictions, it is a method that is both reliable and easy to set up; however, secure ICAP might result in added expense when running your network.

To offset the added expense of running secure ICAP, other alternatives can be considered; however, these alternatives do depend on network topology and platform considerations.


This section discusses three methods to consider when setting up your ICAP server.

- ❑ ["Using Secure ICAP" on page 488](#)
- ❑ ["Using a Crossover Cable" on page 491](#)
- ❑ ["Using a Private Network" on page 493](#)

Using Secure ICAP

Secure ICAP allows you to run ICAP over an encrypted channel. Encrypting the data between the ProxySG appliance and the ICAP server protects the integrity of messages that are sent between the two machines, as it blocks impersonators and prevents a man-in-the-middle attack.

Secure ICAP relies on an SSL device profile that validates a client certificate against a chosen CA Certificate List (CCL) to verify authentication. Secure ICAP presents a server certificate to the ProxySG appliance, after which, the ProxySG appliance must verify the results before a connection is permitted.



The network diagram shows an SG210 appliance connecting to an ICAP server. Because the SG210 has limited physical port capabilities, using secure ICAP is an effective and easy solution to connect the two devices.

Note: The SG210 has only two network interfaces that default as a hardware bridge.

To configure secure ICAP:

The following procedure assumes you are using a ProxyAV appliance as your ICAP server.

1. Ensure that the ProxyAV appliance is set up and configured for Secure ICAP.
 - a. From the ProxyAV appliance console, click on ICAP Settings and verify that the **Secure** check box is enabled.
2. Copy the "Default" SSL certificate that will be imported by the ProxySG appliance. On the ProxyAV appliance, select **Advanced > SSL Certificates**.
 - a. Select **Default** and copy the certificate. You must include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` statements.
3. From the ProxySG appliance's Management Console, copy the ProxyAV appliance's default SSL certificate to the CA Certificate List.
 - a. Select **Configuration > SSL > CA Certificates**.
 - b. Click **Import**. The **Import CA Certificate** dialog displays.

- c. Choose a name for this certificate and enter it in the **CA Cert Name** field, then paste the certificate in the **CA Certificate PEM** panel. You must include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` statements.
 - d. Click **OK**. The dialog closes and you return to the **CA Certificates** tab.
 - e. Click **Apply** to save your settings.
4. Create a CA Certificate List. For more information, see ["Managing CA Certificate Lists"](#) on page 1143.
 - a. From the ProxySG appliance Management Console, select **Configuration > SSL > CA Certificates > CA Certificate Lists**. The **CA Certificates Lists** page displays.
 - b. Click **New**. The **Create CA Certificate List** dialog displays.
 - c. Enter the name of the certificate list in the field provided.
 - d. Locate the certificate that you imported in Step 3, then click **Add >>** to move the certificate to the **Selected** column.
 - e. Click **OK**. The dialog closes and you return to the **CA Certificate Lists** page. The new certificate list is shown in the table.
 - f. Click **Apply** to save your settings.
5. Create an SSL device profile for the ICAP server. For more information, see ["About SSL Device Profiles"](#) on page 1291.
 - a. From the ProxySG appliance Management Console, select **Configuration > SSL > Device Profiles**. The **Profiles** page displays.
 - b. Click **New**. The **Create SSL Device Profile** dialog displays.
 - c. Enter the name of the device profile in the field provided.
 - d. Set **Keyring** to **<None>**.
 - e. Select the **CCL** that you created in Step 4 from the drop-down list.
 - f. Enable **Verify peer** by selecting the check box.
 - g. Click **OK**. The dialog closes and you return to the **Profiles** page.
 - h. Click **Apply** to save your settings.
6. Configure ICAP on the ProxySG appliance.
 - a. From the ProxySG appliance Management Console, select **Configuration > External Services > ICAP > ICAP Services**. The **Services** page displays.
 - b. Click **New**. The **Add list item** dialog displays.
 - c. Enter the name of the ICAP service, then click **OK**. The dialog closes and you return to the **Services** page. The new service is listed in the table.


-
- d. Select the ICAP service you just created, then click **Edit**. The **Edit ICAP Service** *ICAP_Service_Name* dialog displays.
 - e. Enter the Service URL of the ICAP server, for example, `icap://192.0.2.0/avscan`.
 - f. In the **ICAP Service Ports** section, select the check box for **This service supports secure ICAP connections**.
 - g. Set the SSL device profile to the profile that was created in Step 5.
 - h. To set remaining configurations for the ProxySG appliance, see ["Creating an ICAP Service"](#) on page 474.
 - i. Click **OK**. The dialog closes and you return to the **ICAP Services** page.
 - j. Click **Apply** to save your settings.

Using a Crossover Cable

As an alternative to using encryption between a ProxySG appliance and an ICAP server, you can use an Ethernet crossover cable to connect the ProxySG appliance directly to an ICAP server, such as Content Analysis or a ProxyAV appliance, to secure the connection. This solution can be applied to all ProxySG appliances except the SG210.

Note: The SG210 offers only two physical ports, making it difficult to connect to the private network and WAN, while still being able to plug in the Ethernet cable.

The crossover cable provides a physical connection between the ProxySG appliance and the ICAP server. The data is sent unencrypted and no authentication is required. Using a crossover cable as a means to connect the two devices is an effective alternative to using secure ICAP, as long as the connection is in a secure and controlled environment.



Configuring ICAP Using a Crossover Cable

To configure ICAP using a crossover cable:

1. Plug the ProxySG appliance interface 0:0 and ICAP server interface 0 into your network in the usual way and configure appropriate IP addresses, DNS, etc., on each.
2. Define a subnet with two IP addresses (/30), assigning one IP address to the ProxySG appliance and the other IP address to the ICAP server, making sure neither overlaps with any active subnets. This example is using netmask 255.255.255.252 (/30) as the netmask because only two IP addresses are required; however, larger subnets can be used.
3. Set the ProxySG appliance's IP address from this subnet on its interface 1:0, with selected netmask.
4. Set the ICAP server's IP address from this subnet on its interface 1 using the same netmask.
5. Plug the ICAP server's interface 1 into the ProxySG appliance's interface 1:0 with a crossover cable.
6. Create one or more ICAP services on the ProxySG appliance, pointing at the ICAP server's IP address selected above.
 - a. Create an ICAP response service. Select **Configuration > External Services > ICAP > ICAP Services**. The **ICAP Services** page displays.


-
- b. Select **New**. The **Add list item** dialog displays. Enter the ICAP service in the field provided, then click **OK**. The dialog closes and you return to the **ICAP Services** page.
 - c. Select the service you just created, then click **Edit**. The **Edit ICAP Service** dialog displays.
 - d. Enter the Service URL of the ICAP server, for example, `icap://192.0.2.0/avscan`.
 - e. In the **ICAP Service Ports** section, select the check box for **This service supports plain ICAP connections**.
 - f. To set remaining configurations for the ProxySG appliance, see ["Creating an ICAP Service"](#) on page 474.
 - g. Click **OK**. The dialog closes and you return to the **ICAP Services** page.
 - h. Click **Apply** to save your settings.

Using a Private Network

Larger enterprises may require redundancy in the network (for example, multiple ProxySG appliances and/or multiple ICAP servers such as Content Analysis or ProxyAV appliances). Redundant appliances address the limitations of the single ICAP server/single-ProxySG deployment. The ProxySG appliance can load balance Web content scanning between multiple ICAP servers, or designate a sequence of ICAP servers as failover devices should the primary ICAP appliance go offline. Similarly, secondary ProxySG appliances can be configured as failover devices should the primary ProxySG appliance go down and can provide further proxy support in the network.

If you have multiple ProxySG appliances that share multiple ICAP servers, you can configure a private network that is completely separate from other networks within your organization.

The figure shows a ProxySG appliance using multiple ICAP servers interconnected on a private network.



To configure a private network:

1. Plug each ProxySG appliance interface `0:0` and each ICAP server interface `0` into your network in the usual way and configure appropriate IP addresses, DNS, etc., on each.
2. Define a subnet with two IP addresses (`/30`), assigning one IP address to the ProxySG appliance and the other IP address to the ICAP server, making sure neither overlaps with any active subnets. This example is using netmask `255.255.255.252` (`/30`) as the netmask because only two IP addresses are required, however, larger subnets can be used. Make sure you define a subnet large enough to assign addresses to all ProxySG appliances and ICAP servers that are being interconnected.

Symantec recommends that all ProxySG appliances reside on the same subnet as the ICAP servers, even in cases where multiple ProxySG appliances are load balanced with multiple ICAP servers. Although you can put the ICAP server in California and the ProxySG appliance in New York, performance will suffer. For optimal performance, the ICAP server and ProxySG appliance must be physically and logically close to each other; Symantec recommends that the ICAP server be on the next-hop VLAN.

3. Set each ProxySG appliance's IP address from this subnet on its interface `1:0`, with selected netmask.
4. Set each ICAP server's IP address from this subnet on its interface `1` using the same netmask.
5. Plug each ProxySG appliance's interface `1:0` and each ICAP server's interface `1` into the private network switch.

-
6. Create one or more ICAP services on each ProxySG appliance, pointing at the ICAP servers' IP addresses selected above.
 - a. Create an ICAP response service. Select **Configuration > External Services > ICAP > ICAP Services**. The **ICAP Services** page displays.
 - b. Select **New**. The **Add list item** dialog displays. Enter the ICAP service in the field provided, then click **OK**. The dialog closes and you return to the **ICAP Services** page.
 - c. Select the service you just created, then click **Edit**. The **Edit ICAP Service** dialog displays.
 - d. Enter the Service URL of the ICAP server, for example, `icap://192.0.2.0/avscan`.
 - e. In the **ICAP Service Ports** section, select the check box for **This service supports plain ICAP connections**.
 - f. To set remaining configurations for the ProxySG appliance, see ["Creating an ICAP Service"](#) on page 474.
 - g. Click **OK**. The dialog closes and you return to the **Services** page.
 - h. Click **Apply** to save your settings.

Section D: Monitoring ICAP Requests and Sessions

This section discusses the following topics:

- ❑ ["Introduction to ICAP Request Monitoring"](#)
- ❑ ["ICAP Graphs and Statistics"](#) on page 496
- ❑ ["Monitoring ICAP-Enabled Sessions"](#) on page 499

Introduction to ICAP Request Monitoring

After configuring ICAP services, you can monitor the transactions and connections to validate ICAP functionality and analyze ICAP issues. For example, you can determine how many scanning requests were successful versus failed in a certain time period.

- ❑ For displaying tabular statistics and graphing historical ICAP data, use the ICAP statistics page. See ["ICAP Graphs and Statistics"](#) on page 496.
- ❑ For monitoring ICAP-enabled sessions, use the Active Sessions and Errored Sessions pages. See ["Monitoring ICAP-Enabled Sessions"](#) on page 499.

ICAP Graphs and Statistics

You can display a variety of ICAP statistics in bar chart form as well as in a statistical table. The following table defines the ICAP statistics that the ProxySG tracks for each ICAP service and service group.

Table 22–3 ICAP Statistics

Statistic	Definition
Plain Requests	ICAP scanning transactions that are not encrypted
Secure Requests	ICAP scanning transactions that are encrypted and tunneled over SSL
Deferred Requests	ICAP scanning transactions that have been deferred until the full object has been received
Queued Requests	ICAP scanning transactions that are waiting until a connection is available
Successful Requests	ICAP scanning transactions that completed successfully
Failed Requests	ICAP scanning transactions that failed because of a scanning timeout, connection failure, server error, or a variety of other situations
Bytes Sent	Bytes of ICAP data sent to the ICAP service or service group Note: Bytes Sent does not include secure ICAP traffic.
Bytes Received	Bytes of data received from the ICAP service or service group

Table 22–3 ICAP Statistics


Statistic	Definition
Plain Connections	Line of communication between the ProxySG appliance and an ICAP server across which plain ICAP scanning requests are sent Note: This statistic is not tracked for service groups.
Secure Connections	Secure line of communication between the ProxySG appliance and an ICAP server across which encrypted ICAP scanning requests are sent Note: This statistic is not tracked for service groups.

Displaying ICAP Graphs

ICAP graphs can be used as diagnostic and troubleshooting tools. For instance, if the Active Requests graph shows excessive queued ICAP requests on a regular basis, this may indicate the need for a higher capacity ICAP server.

To display an ICAP graph:

1. Select **Statistics > ICAP**. The **ICAP** statistics screen displays.



2. Select whether to graph **Services** or **Service Groups**.
3. From the **Duration** drop-down list, select the time period to graph: **Last Hour**, **Last Day**, **Last Week**, **Last Month**, or **Last Year**.
4. Select the type of graph:
 - Active Requests** — Plain, secure, deferred, and queued active ICAP transactions (sampled once per minute)
 - Connections** — Plain and secure ICAP connections (sampled once per minute)
 - Completed Requests** — Successful and failed completed ICAP transactions


Bytes — Bytes sent to the ICAP service and received from the ICAP service

Each statistic displays as a different color on the stacked bar graph. By default, all relevant statistics are displayed.

5. In the **Name** column in the table beneath the graph, select the service or service group you wish to graph or select the **Totals** row to graph all services or service groups.
6. (Optional) Clear the options next to any statistics that you do not want displayed on the graph.

Additional Information

- ❑ While the ICAP statistics screen is displayed, you can view new graphs by selecting different services, service groups, time periods, or graph types.
- ❑ Graphs automatically refresh every minute. This may be noticeable only on graphs with the Last Hour duration.
- ❑ To see the actual statistics associated with a bar on the graph, hover the mouse pointer anywhere on the bar. A box showing the statistics and total appears at the mouse pointer.



Displaying ICAP Statistical Data

If you are more interested in the *data* than in the graphs, the **ICAP** statistics screen displays this information as well; beneath the graph is a concise table that displays the number of successful and failed requests and number of bytes sent and received for each service or service group during the selected time period. The table also calculates totals for each statistic across all services or service groups.

To display ICAP statistical data:

1. Select **Statistics > ICAP**. The ICAP statistics screen displays.
2. Select whether to display statistics for **Services** or **Service Groups**.
3. From the **Duration** drop-down list, select the time period for the statistics: **Last Hour**, **Last Day**, **Last Week**, **Last Month**, or **Last Year**.

For the time period you selected, the ProxySG displays statistics for individual services as well as totals for all services.

Name ▲	Successful Requests	Failed Requests	Bytes Sent	Bytes Received
icap_req1	3,967	0	1,613,566	
icap_req2	2,399	0	580,314	
icap_resp1	4,388	0	4,271,319,127	
icap_resp2	2,737	1	3,898,808	
Totals :	26,208	1	4,281,824,614	

Monitoring ICAP-Enabled Sessions

For detailed information about active and errored sessions that have ICAP scanning enabled, view the Active Sessions and Errored Sessions pages. You can filter the session list to display only the ICAP-enabled sessions, so that you can easily view the status of each session (transferring, deferred, scanning, completed) and see fine-grained details (such as client IP address, server name, bytes, savings, and protocol).

Additional ICAP filters are available as well. You can also filter by:

- ❑ Type of ICAP service (REQMOD or RESPMOD)
- ❑ Service name
- ❑ ICAP status (for example, display only the deferred sessions)

Additional filters are optional. If you leave all the options set to **Any**, all ICAP sessions are displayed.

Displaying Active ICAP-Enabled Sessions

By default, the **Active Sessions** screen displays all active sessions. When analyzing ICAP functionality, it is helpful to filter the list to display only ICAP-enabled sessions.

To list ICAP-enabled sessions:

1. Select **Statistics > Sessions > Active Sessions > Proxied Sessions**.
2. Select the **ICAP** filter from the **Filter** drop-down list.
3. (Optional) Select the type of ICAP service from the drop-down list: **Any, REQMOD, RESPMOD**.
4. (Optional) Select the service name from the **Service** drop-down list.
5. (Optional) Select the ICAP state from the **Status** drop-down list: **Any, transferring, deferred, scanning, completed**.
6. (Optional) To limit the number of connections to view, select **Display the most recent** and enter a number in the results field. This helps optimize performance when there is a large number of connections.
7. (Optional) To view the current errored proxied sessions, select **Show errored sessions only**.

8. Click **Show**. The Proxied Sessions table displays the ICAP-enabled sessions.

Client	Server	A	S	FW	I	Duration	Client Bytes	Server Bytes	Sav
▶ 10.100.1.16:58988	ads.yimg.com:80					1.7 min	20,801	20,891	
● 10.100.1.16:59849	us.news2.yimg.com:80					1.9 min	3,375	3,420	
● 10.100.1.16:60218	streamerapi.finance.yah...					1.8 min	913	22,511	
● 10.100.1.16:60231	streamerapi.finance.yah...					1.6 min	874	919	

Terminate Session Download

Of particular interest in the Proxied Sessions table is the ICAP (I) column. This column indicates the status of the ICAP-enabled session, with unique icons identifying the status of the connection. [Table 22-4](#) describes each of the icons. For descriptions of the other columns in the table, see ["About the Proxied Sessions Statistics"](#) on page 694.

Table 22-4 ICAP icons

ICAP Icon	Description
(magnifying glass) 	Scanning — ICAP requests are in the process of being scanned
(arrow) 	Transferring — ICAP requests are being transferred to the ICAP server
(clock) 	Deferred — ICAP scanning requests have been deferred until the full object has been received
(check mark) 	Completed — ICAP scanning requests completed successfully
(i) 	Inactive — The ICAP feature is inactive for the session or connection
no icon	Unsupported — ICAP is not supported for the corresponding session or connection

Additional Information

- ❑ Icon Tooltips—When you mouse over an ICAP icon, a tooltip displays details about the ICAP-enabled session:
 - The type of service (REQMOD and/or RESPMOD)
 - The name of the service

- The ICAP state (transferring, deferred, scanning, or completed), for example:

REQMOD Service: icap1 (completed)

- ❑ When the following conditions are met, two ICAP services display for one explicit HTTPS connection:
 - An ICAP service group is used for request modification (REQMOD) and there are more than one ICAP service in the ICAP service group.
 - Explicit HTTPS connections are set by policy to perform ICAP request modification (REQMOD).
 - The ProxySG is configured to intercept these HTTPS connections.
- ❑ When only one type of service is used for a session, the tooltip indicates whether the other type is inactive or unsupported, for example:
RESPMOD Service: inactive

Sorting—If you click the **I** column heading, the sessions are sorted in the following order:

- ❑ Transferring
- ❑ Deferred
- ❑ Scanning
- ❑ Completed
- ❑ Inactive
- ❑ Unsupported

Displaying Errored ICAP-Enabled Sessions

As with active sessions, errored sessions can be filtered to display only ICAP-enabled sessions.

To filter the errored session list to display only ICAP-enabled sessions:

1. Select **Statistics > Sessions > Errored Sessions > Proxied Sessions**.
2. Select the **ICAP** filter from the **Filter** drop-down list.
3. (Optional) Select the type of ICAP service from the drop-down list: **Any, REQMOD, RESPMOD**.
4. (Optional) Select the service name from the **Service** drop-down list.
5. (Optional) Select the ICAP state from the **Status** drop-down list: **Any, transferring, deferred, scanning, completed**.
6. (Optional) To limit the number of sessions to view, select **Display the most recent** and enter a number in the results field. This helps optimize performance when there is a large number of connections.

- Click **Show**. The Proxied Sessions table displays the active and inactive errored ICAP-enabled sessions.

Proxied Sessions

Client ▲	Server	A	S	FW	I	Duration	Client Bytes	Server Bytes	Save
● 10.100.1.16:60219	10.100.1.10:80				✓	0 sec	1,278	5,865	▲
● 10.100.1.16:60231	streamerapi.finance.yah...				✓	6 min	874	1,833	
▶ 10.100.1.16:60932	10.100.1.10:80				✓	29 sec	47,234	52,560	
▶ 10.100.1.16:60936	10.100.1.10:80				✓	16 sec	6,485	6,378	▼

Terminate Session Download Clear

Section E: Creating ICAP Policy

While the ICAP service defines the parameters for setting up the transaction between the ProxySG and the ICAP server, ICAP policy allows you to specify the response or action for each ICAP service or service group that is configured. For example, using policy, you may have a general rule for scanning all incoming responses and set an action to deny content if the scan cannot be completed. You then can create a rule that allows responses from specific business critical sites to be served even if the scan cannot be completed. Or, for super users in your enterprise, you can allow access to password protected archives whose content cannot be scanned.

Policy allows for creating granular rules based on individual users, groups of users, time of day, source, protocol, user agent, content type and other attributes. For example, you can create policy to define an action when a virus is detected, or when an ICAP error or ICAP server failover occurs in your network. Policy can be created using the graphical user interface, the Visual Policy Manager (VPM) or Content Policy Language (CPL).

The following topics are discussed in this section:

- ❑ ["VPM Objects" on page 503](#)
- ❑ ["Example ICAP Scanning Policy" on page 504](#)
- ❑ ["Exempting HTTP Live Streams From Response Modification" on page 509](#)
- ❑ ["Streaming Media Request Modification Note" on page 509](#)
- ❑ ["Using ICAP Error Codes in Policy" on page 509](#)
- ❑ ["Using ICAP Headers in Policy" on page 514](#)
- ❑ ["CPL Notes" on page 516](#)

VPM Objects

The VPM contains the following objects specific to Web content scanning.

Table 22–5 AV Scanning Objects

Object	Layer>Column
<i>Virus Detected</i>	<i>Web Access>Service</i>
<i>ICAP Error Code</i>	<i>Web Access>Service</i>
<i>Return ICAP Feedback</i>	<i>Web Access>Action</i>
<i>Set ICAP Request Service</i>	<i>Web Access>Action</i>
<i>Set ICAP Request Service</i>	<i>Web Content>Action</i>
<i>Set ICAP Response Service</i>	<i>Web Content>Action</i>
<i>Set Malware Scanning</i>	<i>Web Content> Action</i>
<i>ICAP Respmod Response Header</i>	<i>Web Access > Destination</i>
<i>ICAP Reqmod Response Header</i>	<i>Web Access > Source</i>

For information on the VPM and defining policies, refer to *Visual Policy Manager Reference*.

For more information on using CPL, refer to the *Content Policy Language Reference*.

Example ICAP Scanning Policy

The following VPM example demonstrates the implementation of an ICAP policy that performs virus scanning on both client uploads (to prevent propagating a virus) and responses (to prevent the introduction of viruses), and provides failover with backup ICAP services.

For this example:

- ❑ The ProxySG has configured ICAP services. The response service is **avresponse1** and the request service is **avrequest1**.
- ❑ Two backup response services are configured: **avresponse2** and **avresponse3**.
- ❑ The CAS/ProxyAV is the virus scanner and it is configured to serve, (rather than block) password-protected files.
- ❑ A group named IT is configured on the ProxySG.
- ❑ The IT group wants the ability to download password protected files, but deny everyone else from doing the same.

To perform virus scanning, protecting both the server side and the client side:

1. In the VPM, select **Policy > Web Access Layer**. Name the layer **RequestAV**.
2. Right-click the **Action** column; select **Set**. The Set Action Object dialog displays.
3. Click **New**.
4. Select **Set ICAP Request Service**; the Add ICAP Request Service Object dialog displays.

BC Add Request Analysis Service Object

Name: RequestAnalysisService1

5a ☐ Do not perform request analysis

5b ☐ Use the internal request analysis service

☒ Use the following external request analysis services

1 ICAP request services configured on SG

5c Available services: [Empty List] Add >> Move Up Move Down << Remove

Selected failover sequence: request_service

5d Connection Security

☒ If available use secure ICAP connections for encrypted requests

☐ Always use secure ICAP connections

☐ Always use plain ICAP connections

5e If the request analysis service is not available:

☒ Deny the client request (recommended)

☐ Continue without further request processing

OK Cancel Help

5. Configure the request service object:
 - a. Select **Add Request Analysis service**.
 - b. If you will be using the internal Content Analysis service, select **Use the Internal Request Analysis service**. If you're configuring an external Content Analysis or ProxyAV server, continue with the remaining steps.
 - c. If using an external ICAP service, Select the ICAP mode **If available use secure ICAP connections for encrypted responses**. This mode uses plain ICAP for HTTP and FTP traffic and secure ICAP for HTTPS traffic. This is the default mode.

The **Always use secure ICAP connections** mode uses secure ICAP for all traffic (HTTP, HTTPS, FTP). The **Always use plain ICAP connection** mode uses plain ICAP for all traffic (HTTP, HTTPS, FTP).


- d. From the **Available services** field, select the **avrequest1** and click **Add**. This moves the service name to the **Selected failover sequence** field.

- e. Accept the default: **Deny the client request**. This prevents a client from propagating a threat. If a virus is found, the content is not uploaded. For example, a user attempts to post a document that has a virus and is denied.
- f. Click **OK**; click **OK** again to add the object to the rule.

RequestAV							
No.	Source	Destination	Service	Time	Action	Track	Comment
1	Any	Any	Any	Any	 RequestAV	None	

Figure 22–6 Request

6. In the VPM, select **Policy > Add Web Content Layer**. Name the rule **ResponseAV**.
7. Right-click the **Action** column; select **Set**. The Set Action Object dialog displays.
8. Click **New**.
9. Select **Set ICAP Response Service**; the Add ICAP Response Service Object dialog displays.



10. Configure the response service object:
 - a. Select Add Response Analysis Service Use **ICAP response service**.
 - b. Select the ICAP mode, **If available use secure ICAP connections for encrypted requests**.
 - c. Select **avresponse1** and click **Add**.
 - d. Repeat Step b for to add the additional failover services.
 - e. Select **Deny the client request**. This scans the responses for viruses before the object is delivered to the client. If a virus is found, the content is not served.
 - f. Click **OK**; click **OK** again to add the object to the rule.

To log a detected virus:


1. In the VPM, select **Policy > Web Access Layer**. Name the layer **AVErrors**.
2. Right-click the **Service** column; select **Set**. The Set Service Object dialog displays.
 - a. Select **Virus Detected** (static object).
 - b. Click **OK** to add the object to the rule.
3. Right-click the **Action** column. Select **Deny**.
4. Right-click the **Track** column. Select **Set**; the Set Track Object dialog displays.
 - a. Click **New**; select **Event Log**. The Event Log dialog displays.
 - b. In the **Name** field, enter **VirusLog1**.
 - c. From the scroll-list, select **icap_virus_details**, **localtime**, and **client-address**. Click **Insert**.
 - d. Click **OK**; click **OK** again to add the object to the rule.

RequestAV ResponseAV AVErrors							
No.	Source	Destination	Service	Time	Action	Track	Comment
1	Any	Any	 Virus Detected	Any	 Deny	 VirusLog1	

Figure 22–7 The AVErrors rule

To create an exception for IT group:

1. In VPM, select **Policy > Add Web Access Layer**. Name the rule **AVExceptions**.
2. Add the **IT** group object to the **Source** column.
3. Right-click the **Service** column; select **Set**. The Set Service Object dialog displays.
4. Click **New**; select **ICAP Error Code**. The **Add ICAP Error Code** Object displays.



5. Add the error code:
 - a. Select **Selected Errors**.
 - b. From the list of errors, select **Password Protected Archive**; click **Add**.
 - c. Name the object **password_protected**.
 - d. Click **OK**; click **OK** again to add the object to the rule.
6. Right-click the **Action** column and select **Allow**.
7. Click **Add Rule**.
8. In the **Service** column, add the **password_protected** object.
9. Right-click the **Action** column; select **Deny**.

RequestAV ResponseAV AVErrors AVEExceptions							
No.	Source	Destination	Service	Time	Action	Track	Comment
1	cn=IT...	Any	password_protected	Any	Allow	None	
2	Any	Any	password_protected	Any	Deny	None	

After this policy is installed:

- ❑ Malware scanning is performed for client attempts to upload content and content responses to client requests.
- ❑ If malware is detected and there were no scanning process errors, a log entry occurs.
- ❑ As the CAS/ProxyAV is configured to serve password-protected objects, only the IT group can download such files; everyone else is denied.

Exempting HTTP Live Streams From Response Modification

The following CPL examples demonstrate how to exempt HTTP live streams from response modification, as they are not supported by ICAP. The CPL designates user agents that are bypassed.

```
<cache>
url.scheme=http request.header.User-Agent="RealPlayer G2"
response.icap_service(no)
url.scheme=http request.header.User-Agent="(RMA) "
response.icap_service(no)
url.scheme=http request.header.User-Agent="(Winamp) "
response.icap_service(no)
url.scheme=http request.header.User-Agent="(NSPlayer) "
response.icap_service(no)
url.scheme=http request.header.User-Agent="(Windows-Media-Player) "
response.icap_service(no)
url.scheme=http request.header.User-Agent="QuickTime"
response.icap_service(no)
url.scheme=http request.header.User-Agent="(RealMedia Player) "
response.icap_service(no)
```

Streaming Media Request Modification Note

Some HTTP progressive download streaming media transactions are complex enough to disrupt ICAP request modification services. If such behavior is noticed (most common with RealPlayer), implement a workaround policy to bypass the ICAP request modification service for HTTP progressive downloads:

For example:

```
<proxy>
url.scheme=http request.header.User-Agent="(RealMedia Player) "
request.icap_service(no)
url.scheme=http request.header.User-Agent="RMA"
request.icap_service(no)
```

Using ICAP Error Codes in Policy

ICAP error codes are available as objects in policy for the CAS/ProxyAV ICAP server only and are useful for creating policy that is flexible and granular.

ICAP error codes are available in the **Service** column of the **Web Access Layer**. For each error code, an action can be defined in policy. For example, if your default policy is set to deny requests when an ICAP scan cannot be completed, the user will be denied access to the content when a CAS/ProxyAV is unavailable to process requests. To prevent the user from being denied access, you can create policy to allow access to specific sites without ICAP scanning when the ICAP error code is **Server Unavailable**. This policy allows requests to the specified sites without ICAP scanning when the CAS/ProxyAV is unavailable for content scanning.

The following table lists the error codes and their descriptions:

Table 22–6 ICAP Error Codes Available in Policy

ICAP Error Code	VPM Object Name	Description
Errors generated by the CAS/ProxyAV These antivirus scanning options are available on the Antivirus Settings > Scanning Behavior link of the CAS/ProxyAV Management Console.		
Scan timeout	Scan Timeout	Scan operation was abandoned because the file scanning timeout was reached. The default is 800 seconds.
Decode error	Decode Error	Error detected during file decompression/decoding.
Password protected	Password Protected Archive	Archive file could not be scanned because it is password protected.
Insufficient space	Insufficient Space	Indicates that the disk is full.
Max file size exceeded	Maximum File Size Exceeded	Maximum individual file size to be scanned exceeds settings in configuration. The maximum individual file size that can be scanned depends on the RAM and disk size of the ProxyAV appliance model.
Max total size exceeded	Maximum Total Size Exceeded	Maximum total uncompressed file size exceeds settings in configuration. The maximum limit varies by ProxyAV appliance model.
Max total files exceeded	Maximum Total Files Exceeded	Maximum total files in an archive exceeds settings in configuration. The maximum is 100,000.
Max archive layers exceeded	Maximum Archive Layers Exceeded	Maximum number of layers in a nested archive exceeds settings in configuration. The maximum by vendor is: • Panda: 30 • McAfee: 300 • All others: 100.
File type blocked	File Type Blocked	Blocked a file type as configured on the ICAP server settings.
File extension blocked	File Extension Blocked	Blocked a file extension as configured on the ICAP server settings.
Antivirus load failure	Anti-virus Load Failure	Unable to load antivirus engine on the ICAP server.
Antivirus license expired	Anti-virus License Expired	Antivirus license expired.
Antivirus engine error	Anti-virus Engine Failure	Antivirus engine error.

ICAP Error Code	VPM Object Name	Description
Error messages generated by the ProxySG		
ICAP connection mode not supported	ICAP Connection Mode not Supported	ICAP server does not support the configured connection mode. For example, plain ICAP is required but server supports only secure ICAP and vice versa.
ICAP security error	ICAP Security Error	(Secure ICAP error) Unable to establish a secure connection to the ICAP server. This could be because the SSL device profile is not enabled or is corrupt.
Connection failure	Connection Error	Unable to connect to the ICAP server—applies to connection refused, connection timed out, or any other error when connecting. It would also apply if the connection dropped unexpectedly while sending a request or reading a response.
Request timeout	Request Timeout	Request timed out because no response was received from the ICAP server within the configured connection timeout, although the connection to the server is healthy. The default connection timeout is 70 seconds.
Internal error	Internal Error	Description varies and implies an internal processing error on the ProxySG.
Server error	Server Error	Displayed when the ProxySG receives a 4xx or 5xx error from the ICAP server that does not contain the error code and error details.
Server Unavailable	Server Unavailable	Unable to process an ICAP request because the ICAP server in the service/service group is unhealthy.

Example of Using an ICAP Error Code in Policy


The following example illustrates how to create policy to serve or deny content when the Decode/Decompression ICAP error code is triggered.

When scanning a file for viruses, the scan engine might return a decompression error. A decompression error is triggered by the scan engine when it interprets an invalid form of file compression; the ProxyAV appliance has no control over this error.


Because the scan engine perceives this error as a security threat, you can create policy to block these files for most users but serve the unscanned content for select user groups in your enterprise.

To create a policy for the Decode/Decompression ICAP error code:


1. Launch the VPM and select the **Web Access Layer**.
2. Add an **ICAP Error Code** object:
 - a. In the **Service** column, right click and select **Set**. The **Set Service Object** dialog displays.
 - b. Click **New** and select **ICAP Error Code**. The **Add ICAP Error Code** dialog displays.



- c. Click **Selected Errors** and select the **Decode error** from the list of available errors.
 - d. Click **OK** to save your changes and exit all open dialogs.
3. In the **Source** column, right click and select **Set**. The **Set Source Object** dialog displays.



- a. Select **Group**. The **Add Group Object** dialog displays.
- b. Add the group and authentication realm for the users with access to unscanned content. Click **OK** to save your changes and exit the **Add Group Object** dialog.
4. In the **Action** column, right-click set the action to **Allow**. Now you have rule in the **Web Access Layer** that allows the group access to content when the ProxySG receives the ICAP decode error.
5. Add another rule in the **Web Access Layer** to deny access to unscanned content to all other users in the network.



- a. In the **Service** column, right click and select **Set**. The **Set Service Object** dialog displays.
 - b. Select the ICAP error code service object that you created in Step 1 from the list. Click **OK**.
 - c. In the **Action** column, right click and set action to **Deny**.
6. Click **Install Policy** to install the policy.

Using ICAP Headers in Policy

When you enable ICAP, traffic through the ProxySG appliance is sent to ProxyAV for scanning. ProxyAV scans the content and may return useful information in ICAP headers to the ProxySG appliance. As an administrator, you can make policy decisions based on the ICAP headers, which can contain information about the scanned files such as virus information, content categorization, and threat levels.

To use ICAP REQMOD headers in policy:

1. Make sure that the appropriate ICAP request or response service has been created on the ProxySG appliance. See ["Creating an ICAP Service"](#) on page 474.
2. Launch the VPM and select the **Web Access Layer**.
3. To inspect ICAP request headers: In the Source column, right click and select **Set**.
To inspect ICAP response headers: In the Destination column, right click and select **Set**.
4. Add a new **ICAP Response Header** object:


- a. Select **New > ICAP Reqmod Response Header** or **ICAP Respmo Response Header**. The object dialog appears.
- b. In the **Name** field, enter a custom name or accept the default.
- c. From the **Header Name** menu, specify the name of the header to inspect.
Before you add a header name, the menu is empty. Any header names you add are saved in the list so you can select them in the future.
- d. In the **Regex** field, enter the pattern to match.
- e. Click **OK**.

Example of Using ICAP Header in Policy

The following policy denies executable (EXE) files. Through an ICAP scan, ProxyAV can determine the apparent data type of the object and return this information in an ICAP header.

To create a policy for ICAP REQMOD response headers:

1. Make sure that the ICAP request service has been created on the ProxySG appliance. See ["Creating an ICAP Service"](#) on page 474.
2. Launch the VPM and select the **Web Access Layer**.
3. In the Source column, right click and select **Set**.
4. Add a new **ICAP REQMOD Response Header** object:
 - a. Select **New > ICAP Reqmod Response Header**. The object dialog appears.



- b. In the **Name** field, enter a custom name or accept the default.
 - c. From the **Header Name** menu, enter "X-Apparent-Data-Types".
 - d. In the **Regex** field, enter "EXE". This refers to the signature for EXE files.
 - e. Click **OK**.
5. In the **Action** column, right click and select **Deny**.
6. Click **Install Policy** to install the policy.

CPL Notes

The following CPL properties are available to manage ICAP services:

- `request.icap_service()` for request modification
- `response.icap_service()` for response modification
- ❑ If policy specifies that an ICAP service is to be used, but the service is not available, the default behavior is to fail closed—that is, deny the request or response. The following CPL allows the serving of objects without ICAP processing if the server is down.

```
request.icap_service(service_name, fail_open)
response.icap_service(service_name, fail_open)
```

When the ICAP service is restored, these objects are scanned and served from the cache if they are requested again.

Note: Blue Coat recommends this CPL to be used for internal sites; use with caution.

- ❑ To provide an exception to a general rule, the following CPL negates ICAP processing:

```
request.icap_service(no)
response.icap_service(no)
```

- ❑ When configuring the secure ICAP feature, the following CPL is used:

Note: This CPL allow the user to configure the `secure_connection` separately for each service in failover sequence.

```
request.icap_service.secure_connection(option)
response.icap_service.secure_connection(option)
request.icap_service.secure_connection.service_name(option)
response.icap_service.secure_connection.service_name(option)
request.icap_service.secure_connection
    [service__0,service_1,...,service_N-1](option)
response.icap_service.secure_connection
    [service__0,service_1,..., service_N-1](option)
```

where `option` is `yes`, `no` or `auto`. The default option is `auto`.

- `yes`— This option means that secure ICAP is used for all traffic (HTTP and HTTPS).
- `no`— This option means that plain ICAP is used for all traffic (HTTP and HTTPS).
- `auto`— This option (default) means that plain ICAP is used for HTTP traffic and secure ICAP is used for HTTPS traffic.

Section F: Managing Virus Scanning

You might need to perform additional ProxySG maintenance concerning virus scanning, particularly for updates to the virus definition on the ICAP virus scanning server.

This section describes the following topics:

- ❑ ["Using Object-Specific Scan Levels"](#) on page 517
- ❑ ["Improving Virus Scanning Performance"](#) on page 517
- ❑ ["Updating the ICAP Server"](#) on page 518
- ❑ ["Replacing the ICAP Server"](#) on page 518
- ❑ ["Configuring Logging for the ICAP Server"](#) on page 518

For information on configuring in-path threat protection and content scanning using the ProxySG and the CAS/ProxyAV, see ["Configuring Threat Protection"](#).

Advanced Configurations

This section summarizes more-advanced configurations between the ProxySG and multiple ICAP servers. These brief examples provide objectives and suggest ways of supporting the configuration.

Using Object-Specific Scan Levels

You can specify different scanning levels for different types of objects, or for objects from different sources.

This requires a service group of ICAP servers, with each server configured to provide the same level of scanning. For more information, see [Chapter 23: "Configuring Service Groups"](#).

Improving Virus Scanning Performance

You can overcome request-handling limitations of ICAP servers. Generally, ProxySGs can handle many times the volume of simultaneous user requests that ICAP servers can handle.

This requires multiple ICAP servers to obtain a reasonable performance gain. On the ProxySG, define policy rules that partition requests among the servers. If you are going to direct requests to individual servers based on rules, configure in rule conditions that only use the URL. Note that you can increase the scale by using a service group, rather than use rules to partition requests among servers. For more information on using multiple ICAP servers, see [Chapter 23: "Configuring Service Groups"](#). For more information about defining policies, refer to the Managing Policy Files chapter in *Visual Policy Manager Reference*, as well as the *Command Line Interface Reference*.

When the virus definitions are updated, the ProxySG stores a signature. This signature consists of the server name plus a virus definition version. If either of these changes, the ProxySG checks to see if the object is up to date, and then rescans it. If two requests for the same object are directed to different servers, then the scanning signature changes and the object is rescanned.

Updating the ICAP Server

If there is a problem with the integration between the ProxySG and a supported ICAP server after a version update of the server, you might need to configure the preview size the appliance uses. For information, see ["Creating an ICAP Service"](#) on page 474.

Replacing the ICAP Server

If you replace an ICAP server with another supported ICAP server, reconfigure the ICAP service on the ProxySG. see ["Creating an ICAP Service"](#).

Configuring Logging for the ICAP Server

The ProxySG provides access log support for Symantec and Finjan ICAP 1.0 server actions (**Management > Access Logging**). The following sections describe access logging behavior for the various supported ICAP servers.

Symantec AntiVirus Scan Engine 4.0

When this Symantec server performs a scan, identifies a problem (for example, a virus), and performs a content transformation, the action is logged. For example:

```
"virus-id: Type=number; Resolution=[0 | 1 | 2]; Threat=name;"
```

where:

- *Type=number* specifies the numeric code for the virus.
- *Resolution=* specifies an integer value that indicates what action was taken to fix the file. Zero (0) defines the file is unrepairable, one (1) specifies that the file was repaired, and two (2) specifies that the file was deleted.
- *Threat=* specifies the name of the virus.

Finjan SurfinGate 7.0

When this Finjan ICAP server performs a scan, identifies a problem (for example, a virus), and performs a content transformation, the action is logged. For example:

```
"virus-id: name, response-info: Blocked, response-desc: virus_name was detected"
```

Finjan ICAP servers also log occurrences malicious mobile code.

Note: The access log string cannot exceed 256 characters. If the header name or value extends the length over the limit, then that string does not get logged. For example, if the `x-virus-id` header value is 260 characters, the access log displays `"x-virus-id:"` with no value because the value is too long to display. Also, if the access log string is already 250 characters and the ProxySG attempts to append a `"Malicious-Mobile-Type:"` string, the string is not appended

Access log entries might vary depending upon the type of ICAP scan performed and the custom log formats. For information about default and custom access log formats, see ["Creating Custom Access Log Formats"](#) on page 649.

Chapter 23: Configuring Service Groups

This chapter describes how to create and manage external ICAP service groups. In high-traffic network environments, a service group accelerates response time by performing a higher volume of scanning.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ "About Service Groups" on page 521
- ❑ "Creating a Service Group" on page 523
- ❑ "Deleting a Service Group or Group Entry" on page 526
- ❑ "Displaying External Service and Group Information" on page 526

About Service Groups

A ProxySG ICAP *service* is a named entity that identifies the ICAP server, the ICAP method, and the supported number of connections. A *service group* is a named set of ICAP services. You will need to create service groups when you are using multiple ICAP servers to process a large volume of scanning requests.

Figure 23–1 shows a service group of three Blue Coat AV ICAP servers.




Figure 23–1 ICAP Service Group

To help distribute and balance the load of scanning requests when the ProxySG is forwarding requests to multiple services within a service group, the ProxySG uses an intelligent load balancing algorithm. When deciding which service in the service group to send a scanning request, this algorithm takes into consideration the following factors:

- ❑ Number of requests that are in a *waiting* state on each service (a request is in this state when it has been sent to the service but the response hasn't been received)
- ❑ Number of unused connections available on each service (calculated by subtracting the number of active transactions from the connection maximum on the server)
- ❑ The user-assigned weight given to each server (see "Weighting" below)

Weighting

Weighting determines what proportion of the load one server bears relative to the others when transactions are waiting to be scanned. (The waiting transactions are typically large file downloads.) If all servers have either the default weight (1) or the same weight, each share an equal proportion of the load when transactions are waiting. If one server has weight 25 and all other servers have weight 50, the 25-weight server processes half as much as any other server.

Before configuring weights, consider the capacity of each server. The processing capacity of the server hardware in relationship to other servers (for example, the number and performance of CPUs or the number of network interface cards) could affect assigned weight of a ICAP server.

Having appropriate weights assigned to your services is critical when all servers in a service group have waiting transactions. As servers reach their capacity, proper weighting is important because requests are queued according to weight.

One technique for determining weight assignments is to start out by setting equal weights to each service in a group; then, after several thousand requests, make note of how many requests were handled by each service. For example, suppose there are two services in a group: Service A handled 1212 requests, Service B handled 2323. These numbers imply that the second service is twice as powerful as the first. So, the weights would be 1 for Service A and 2 for Service B.

Setting the weight value to **0** (zero) disables weighted load balancing for the ICAP service. Therefore, if one ICAP server of a two-server group has a weight value of **1** and the second a weight value of **0**, should the first server go down, a communication error results because the second server cannot process the request.

Load Balancing

When load balancing between services, how does the ProxySG decide which ICAP service to send a scanning request to? For each service, it calculates an index by dividing the number of waiting transactions by the server weight (think of this as wait/weight). The ICAP service with the lowest index value handles the new

ICAP action, assuming that the service has an available connection to use. If it does not, it sends the request to the service with the next lowest index value that has a free connection.

Note: If there are no transactions waiting, load balancing using the assigned weights does not take effect.

Load will be distributed among services proportionally according to their configured weights until the maximum connection limit is reached on all services.

Example 1

Service A and B are in the same service group.

- ❑ Service A can handle up to 50 connections, is assigned a weight of 1, has 17 active transactions, with 5 transactions in the waiting state. The index is calculated by dividing the wait by the weight: $5/1 = 5$.
- ❑ Service B can handle up to 100 connections, is assigned a weight of 2, has 17 active connections, with 15 waiting transactions. The index is $15/2 = 7.5$.

To which service will the ProxySG assign the next ICAP action? Service A because it has a lower index.

Example 2

Service C and D are in the same service group.

- ❑ Service C can handle up to 5 connections, is assigned a weight of 1, has 5 active transactions, with 1 transaction in the waiting state. The index is $1/1=1$.
- ❑ Service D can handle up to 10 connections, is assigned a weight of 1, has 7 active transactions, with 5 waiting transactions. The index is $5/1=5$.


To which service will the ProxySG assign the next ICAP action? Although Service C has a lower index than Service D, it does not have any available connections; therefore, the ProxySG assigns the next ICAP action to Service D which has several free connections.

Creating a Service Group


Create the service group and add the relevant ICAP services to the group. Services within group must be the same type (ICAP).

To configure a service group:

1. Select the **Configuration > External Services> Service-Groups** tab.



2. Add a new group:
 - a. Click **New**; the Add List Item dialog appears.
 - b. In the **Add Service Group** field, enter an alphanumeric name. This example creates a group called **ICAP_Response**.
 - c. Click **OK**.
3. Highlight the new service group name and click **Edit**; the Edit Service Group dialog appears.
4. Select existing services:
 - a. Click **New**; the Add Service Group Entry dialog appears.
 - b. From the list of existing services, select the ones to add to this group. Hold the Control or Shift key to select multiple services.
 - c. Click **OK** to add the selected services to group.



5. Assign weights to services:
 - a. Select a service and click **Edit**; the Edit Service Group Entry weight dialog appears.
 - b. In the **Entry Weight** field, assign a weight value. The valid range is 0-255. For conceptual information about service weighting, see ["Weighting"](#)
 - c. Repeat steps a and b for other services, as required.
 - d. Click **OK** to close the dialog.
 - e. Click **OK** again to close the Edit Service Group Entry dialog
6. Click **Apply**.

When instructed by created policies, the ProxySG sends ICAP response modification requests to ICAP servers in the service group. The load carried by each service in the group is determined by the weight values.

See Also

["About Service Groups"](#) on page 521

["Deleting a Service Group or Group Entry"](#) on page 526

["Displaying External Service and Group Information"](#) on page 526

Deleting a Service Group or Group Entry

You can delete the configuration for an entire service group from the ProxySG, or you can delete individual entries from a service group.

Note: A service or service group used in a ProxySG policy (that is, if a policy rule uses the entry) cannot be deleted; it must first be removed from the policy.

To delete a service group:

1. Select **Configuration > External Services > Service-Groups**.
2. Select the service group to be deleted.
3. Click **Delete**; click **OK** to confirm.
4. Click **Apply**.

To delete a service group entry:

1. Select **Configuration > External Services > Service-Groups**.
2. Select the service group to be modified.
3. Click **Edit**.
4. Select the service entry to be deleted; click **Delete**.
5. Click **OK**.
6. Click **Apply**.

Displaying External Service and Group Information

After configuring a service group, you can display aggregate service group (and other External Services) information.

To display information about all Content Analysis services and groups:

At the (config) command prompt, enter the following commands:

```
SGOS# (config) external-services
SGOS# (config external-services) view
```

Individual service information is displayed first, followed by service group information. For example:

```
; External Services

ICAP-Version:          1.0
URL:                   icap://10.9.59.100/
Plain-ICAP-enabled:    yes
Plain-ICAP-port:       1344
Secure-ICAP-enabled:   no
Secure-ICAP-port:      none
Ssl-device-profile:    none
Max-conn:              25
Timeout (secs):        70
```

```
Defer-threshold:      80%
Notification:         virus-detected
Use ICAP Vendor's virus page:  disabled
Event-log:            connection-failure
Methods:              RESPMOD
Preview-size:         0
Send:                 nothing
ISTag:
Last-ISTag-change:    never
```


Chapter 24: Managing Streaming Media

This chapter describes how to manage streaming content on the enterprise network through the ProxySG streaming proxies.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ [Section A: "Concepts: Streaming Media"](#) on page 530—Explain general streaming concepts and terminology, as well as those specific to the ProxySG streaming solution.
- ❑ [Section B: "Configuring Streaming Media"](#) on page 550—Provides procedures for configuring the ProxySG to manage streaming media applications and bandwidth.
- ❑ [Section C: "Additional Windows Media Configuration Tasks"](#) on page 565—Provides additional procedures for configuring Windows Media.
- ❑ [Section D: "Configuring Windows Media Player"](#) on page 576—Explains how to configure the Windows Media client and describes associated inter activities and access log conventions.
- ❑ [Section E: "Configuring RealPlayer"](#) on page 579—Explains how to configure the Real Media client.
- ❑ [Section F: "Configuring QuickTime Player"](#) on page 583—Describes how to configure the QuickTime client.
- ❑ [Section G: "Using the Flash Streaming Proxy"](#) on page 584—Describes how to configure the ProxySG appliance to manage Flash streaming media applications.
- ❑ [Section H: "Supported Streaming Media Clients and Protocols"](#) on page 592—Describes the vendor-specific streaming protocols supported by the ProxySG.

Section A: Concepts: Streaming Media

This section contains the following topics:


- ❑ ["How the ProxySG Accelerates and Controls Media Streaming"](#) on page 530
- ❑ ["What is Streaming Media?"](#) on page 531
- ❑ ["Streaming Media and Bandwidth"](#) on page 532
- ❑ ["About the Flash Streaming Proxy"](#) on page 532
- ❑ ["About HTTP-Based Streaming"](#) on page 533
- ❑ ["About Windows Media"](#) on page 534
- ❑ ["About Processing Streaming Media Content"](#) on page 538
- ❑ ["IPv6 Support"](#) on page 546
- ❑ ["About Streaming Media Authentication"](#) on page 547

How the ProxySG Accelerates and Controls Media Streaming


The ProxySG streaming media proxies allow you to monitor, control, limit, or even block streaming media traffic on your network. Using the ProxySG for streaming delivery improves the quality of streaming media, reducing artifacts such as frozen playback, and dropped frames or packets. It supports the most popular streaming media clients: Windows Media, Real Media, QuickTime, and Flash.

The ProxySG supports a variety of acceleration, control, and visibility features for streaming media. It provides acceleration features such as live splitting, video-on-demand caching, content pre-population, and multicasting. It also offers control and visibility features such as fine-grained policy control that includes authentication, bandwidth limiting, access logging, and limiting the maximum user connections. The appliance's ability to identify individual users also enables the company to track which employees have watched required videos.

For example, the ProxySG's *pre-population* process can deliver on-demand videos to branch offices during off-hours and save them for future viewing. ProxySG appliances can also cache or save video requested from the headquarters location by a user in a branch office and store it locally for use by subsequent viewers. The diagram below illustrates the process of video caching on the ProxySG.



In the case of live video broadcasts, ProxySG appliances can take a single stream of video and then split it locally into enough streams to serve all local viewers; this is called *live splitting*.



What is Streaming Media?

Streaming media is a term used to describe media files that are served in discrete paced individual packets rather than in bulk, playing while they are being transmitted over the network to the media player on the client computer. In contrast, conventional Web files, which are downloaded through a file transfer, must be downloaded entirely before the user can view them. Commonly requested types of streaming media are video and audio. Streaming media also includes interactive media, cartoon-like animations, panoramic data, and more.

Live versus On-Demand Streaming Media

Streaming media is delivered in the following ways:

- ❑ **Live media streams** Live media streams occur in real time, like the news program that you watch on your television set. Some organizations record a live media stream and then broadcast the media stream to their employees or customers at a specified time. All users who have requested the media stream see the same media stream at the same time. Users are not able to rewind or fast-forward the media stream.
- ❑ **On-demand (previously-recorded) media streams** Users can request these on-demand media streams at a time most convenient to them. Users can pause the media, seek to a different position, rewind, and fast-forward on-demand media streams. On-demand streaming content is commonly referred to as VOD (video-on-demand).

The ProxySG supports both of these types of streaming media.

Streaming Media and Bandwidth

Video, audio, and other streaming media use a considerable amount of bandwidth—much more than the amount of bandwidth needed for Web and news traffic. For example, a media stream could require 10 KB each second, whereas a Web page that the user views for 10 seconds could require 10 KB.

In the typical streaming server-client model, the streaming server sends a separate copy of the media stream to each client that requested the same unique stream. Because streaming media uses a considerable amount of bandwidth, delivering multiple copies of the same media data between the streaming server and the clients can cause significant network and server congestion. The more clients that request the same media stream, the more bandwidth is used.

Planning for efficient bandwidth use is important for streaming media because bandwidth use has a direct correspondence to the quality of the media streams that are delivered to the clients. If your network is congested, your users are likely to experience problems such as jagged video, patchy audio, and unsynchronized video and audio as packets are dropped or arrive late. Conversely, the more bandwidth that is available, the better the quality of media streams.

The ProxySG has several methods for allocating bandwidth to streaming media traffic. See "[Limiting Bandwidth](#)" on page 540.

About the Flash Streaming Proxy

The Flash streaming proxy requires the Flash license. Under a valid trial, demo, or perpetual license, all featured supported by the Flash streaming proxy are enabled. If the license is expired or not installed, the Flash streaming proxy will not accept HTTP-handoff from the HTTP proxy; RTMP traffic tunneled through HTTP proxy using the RTMPT protocol will be handled entirely by the HTTP proxy. Also, if the RTMP proxy listener is set to intercept, those connections are denied.

The Flash proxy provides bandwidth usage optimization for two types of Flash traffic:

- ❑ Live streaming—The ProxySG appliance fetches the live Flash stream *once* from the OCS and serves it to *all* users behind the appliance.
- ❑ Video-on-demand—As Flash clients stream pre-recorded content from the OCS through the ProxySG, the content is cached on the appliance. After content gets cached on the ProxySG, subsequent requests for the cached portions are served from the appliance; uncached portions are fetched from the OCS.

The proxy accelerates plain and encrypted RTMP traffic, both when sent over TCP and when it is tunneled over HTTP. However, the Flash streaming proxy does not support bandwidth limits, or bandwidth management for any RTMP-based protocol, such as RTMP, RTMPT, RTMPE, or RTMPTE.

For additional information, see ["Using the Flash Streaming Proxy"](#) on page 584.

About HTTP-Based Streaming

HTTP-based streaming is an emerging delivery mechanism. Streaming content is encoded at varying bit rates and then fragmented into discrete chunks. The client typically receives a manifest file of the available bit rates and fragments, and can dynamically adapt its request for the next chunk based on client resources (such as CPU) and network conditions (such as bandwidth and congestion).

Limitation

Active Sessions will report HTTP-based streaming protocol information (such as Apple HLS, Adobe HDS, and Microsoft Smooth) as either HTTP or as the appropriate streaming type (**ms_smooth**, **apple_hls**, **adobe_hds**), depending on the types of requests on the connection at any given moment. This is due simply to the nature of the protocols.

About Microsoft Smooth Streaming

One example of HTTP-based streaming is Microsoft's *Smooth Streaming*, which enables adaptive streaming of on-demand and live media over HTTP to clients, such as Silverlight. By dynamically monitoring available bandwidth and video rendering performance, Smooth Streaming optimizes content playback by switching video quality in real-time. For example, users with high bandwidth connections and the latest computers can experience full HD 1080p quality streaming, while users with lower bandwidth or older computers receive a stream that works better for their capabilities.

Smooth Streaming delivers short fragments of video and verifies that each was played back at the expected quality level. If one fragment doesn't play with the expected quality, the next fragment is delivered at a lower quality level. Or, if more bandwidth becomes available, the quality of subsequent fragments will be at a higher level.

When the ProxySG identifies Smooth Streaming over HTTP traffic, the HTTP proxy hands it over to the MS Smooth proxy for processing and reporting. Note that the ProxySG tracks Smooth Streaming traffic separately from other HTTP

traffic and is shown as using the MS Smooth proxy in the Active Sessions, Traffic Mix, and Traffic History reports. You can also create policy to deny or allow streaming requests based on whether the streaming client is Microsoft Smooth Streaming over HTTP.

Note that no additional license is required for Smooth Streaming support.

For additional information, see ["Configuring the HTTP Streaming Proxy"](#) on page 550.

About Adobe HDS

Adobe HDS breaks a video/audio stream into fragments a few seconds long. The files include a manifest (or index) file, which ensures playback in the proper order, and which adapts for quality. The files are in the `.f4m` format.

The ProxySG tracks the HTTP requests carrying Adobe HDS traffic, and presents related information in Active Sessions and in Traffic Mix and Traffic History under Traffic Details in Statistics. The data is reported as Adobe HDS.

No additional license is required for Adobe HDS support.

For additional information, see ["Configuring the HTTP Streaming Proxy"](#) on page 550.

About Apple HLS

Apple HLS (HTTP Live Streaming), developed for iOS and Apple TV devices, is an adaptive streaming technology which breaks a video/audio stream into small fragments, controlled by a "playlist" (or manifest) file (`.m3u8`), which is downloaded at the start of the streaming session. The playlist includes details about the presentation, such as its encryption, supported data rates, and maximum fragment duration, etc. The end tag is not present for live streams, so the client player must periodically re-fetch the playlist. The coded files are distributed in a MPEG-2 transport stream with a `.ts` extension.

The ProxySG tracks the HTTP requests carrying Apple HLS traffic, and presents the related information in Active Sessions and in Traffic Mix and Traffic History under Traffic Details in Statistics. The data is reported as Apple HLS.

No additional license is required for Apple HLS support.

For additional information, see ["Configuring the HTTP Streaming Proxy"](#) on page 550.

About Windows Media

For heightened security and control, some enterprises prefer network environments that restrict Web traffic access (gateway connections) to port 80. Furthermore, beginning with Windows Media Player (WMP) version 11, WMP clients do not use the Microsoft Media Services (MMS) protocol—opting instead for traffic over HTTP and the Real Time Streaming Protocol (RTSP).

Windows Media (WM) streaming over HTTP differs from downloading Windows Media objects over HTTP, which can be stored on any Web server. Streaming content, however, must be hosted on Windows Media Servers that allow the streaming of content over port 80.

SGOS offers unified support for WM content delivered over RTSP and HTTP. The ProxySG appliance's HTTP proxy hands off Windows Media Player HTTP streaming requests to the Windows Media HTTP Module, which itself is a component of the Windows Media RTSP Proxy.

The ProxySG supports the caching of WM content over the RTSP and HTTP protocols. The ProxySG uses the same object cache, which means the content can be served over RTSP and HTTP protocols. WM-HTTP and WM-RTSP both share the same cache.

Live splitting is also supported over both protocols, where all RTSP clients are served by an RTSP splitter and all HTTP clients are served by a separate HTTP splitter, involving two separate live streams to the server, one each for RTSP and HTTP.

Windows Media Deployment

In a Gateway Proxy deployment, the ProxySG supports the caching and splitting of WM content over the RTSP and HTTP protocols. In addition, there are streaming-specific acceleration and policy checks for WM HTTP streaming traffic.

In a Reverse Proxy deployment, the ProxySG can function as a Windows Media server, with WM content delivered over the RTSP and HTTP protocols.

As a Content Delivery Network (CDN) node, the ProxySG supports a shared cache for pre-populated content for delivery over RTSP, RTMP, or HTTP protocols.

Deployment action: Windows Media clients must be configured to enable the HTTP protocol to stream the WM content using HTTP protocol. Similarly, WM clients must be configured to enable RTSP/TCP, and/or RTSP/UDP protocols to stream WM content using RTSP protocol.

Supported Streaming Features

The following table describes the supported Windows Media streaming features.

Live Support

Table 24–1 Windows Media live streaming feature support

Feature	Live Support
Multi-Bit Rate and Thinning	Yes
UDP Retransmission	No
Server-Side Playlists	Yes
Stream Change	Yes

Table 24–1 Windows Media live streaming feature support

Feature	Live Support
Splitting Server-Authenticated Data	Yes
Splitting Proxy-Authenticated Data	Yes

On-Demand Support

Table 24–2 Windows Media on-demand streaming feature support

Feature	On-Demand Support
Multi-Bit Rate and Thinning	Yes
Fast Forward and Rewind	No Caching
Fast Streaming	Yes
UDP Retransmission	No
Server-Side Playlists	No Caching
Stream Change	No
Caching Server-Authenticated Data	Yes
Caching Proxy-Authenticated Data	Yes
Adherence to RTSP Cache Directives	Yes
Partial File Caching	Yes
File Invalidation/Freshness checking for Cached Files	Yes

Multicast Support

Table 24–3 Windows Media multicast UDP streaming feature support

Feature	Multicast
Multi-Bit Rate and Thinning	Yes
Server-Side Playlists	No
Stream Change	No
Multicasting Server-Authenticated Data	No
Multicasting Proxy-Authenticated Data	No

Other Supported Features

The Windows Media streaming feature also supports the following features:

- ❑ Access logging for unicast clients

- ❑ Summary statistics in the Management Console
- ❑ Detailed statistics
- ❑ Forwarding of client streaming logs to origin servers.

Supported VPM Properties and Actions

Windows Media supports the following policy properties and actions:

- ❑ `allow, deny, force_deny`
- ❑ `access_server(yes|no)`. Forces the ProxySG to deliver content only from the cache. Requests for live streams are denied.
- ❑ `authenticate(realm)`
- ❑ `forward(alias_list|no)`
- ❑ `forward.fail_open(yes|no)`
- ❑ `reflect_ip(auto|no|client|vip|<ip address>)`
- ❑ `bypass_cache(yes|no)`. Forces the ProxySG to deliver content in pass-through mode.
- ❑ `limit_bandwidth()`
- ❑ `rewrite()`. One-way URL rewrite of server-side URLs is supported.

Windows Media also supports the following streaming-relevant properties:

- ❑ `max_bitrate(bitrate|no)`. Sets the maximum bit rate that can be served to the client. (This property does not apply to the bit rate consumed on the gateway connection.) If the bit rate of a client-side session exceeds the maximum bit rate set by policy, that client session is denied.
- ❑ `force_cache(yes|no)`. Causes the ProxySG to ignore cache directives and cache VOD content while serving it to clients.
- ❑ `streaming.fast_cache(yes|no)`. Disables the ability of the WM client to request fast-caching of streaming content from the streaming server.

Note: Windows Media does not support policy-based streaming transport selection.

Bandwidth Management

Windows Media supports bandwidth management for both client-side and gateway-side streaming traffic. Bandwidth limits are also supported for pass-through streams. See "[Limiting Bandwidth](#)" on page 540 for more information.

About Processing Streaming Media Content

The following sections describe how the ProxySG processes, stores, and serves streaming media requests. Using the ProxySG for streaming delivery minimizes bandwidth use by allowing the ProxySG to handle the broadcast and allows for policy enforcement over streaming use. The delivery method depends on whether the content is live or video-on-demand.

Delivery Methods

The ProxySG supports the following streaming delivery methods:

- ❑ **Unicast**—A one-to-one transmission, where each client connects individually to the source, and a separate copy of data is delivered from the source to each client that requests it. Unicast supports both TCP- and UDP-based protocols. The majority of streaming media traffic on the Internet is unicast.
- ❑ **Multicast**—Allows efficient delivery of streaming content to a large number of users. Multicast enables hundreds or thousands of clients to play a single stream, thus minimizing bandwidth use.

The following table provides a high-level comparison of unicast and multicast transmission.

Table 24–4 Unicast vs. Multicast

Element	Unicast	Multicast
Connections	One-to-one transmission	One-to-many transmission
Transport	TCP, UDP, HTTP	IP multicast channel
Type of stream	Video-on-demand or live streams	Live streams only
Device requirement	The network devices use unicast.	The network devices must support multicast (not all do).

Serving Content: Live Unicast

A live broadcast can either be truly live or can be of pre-recorded content. A common example is a company president making a speech to all employees.

A ProxySG can serve many clients through one unicast connection by receiving the content from the origin content server (OCS) and then splitting that stream to the clients that request it. This method saves server-side bandwidth and reduces the server load.

Note that you cannot pause or rewind live broadcasts.

Serving Content: Video-on-Demand Unicast

With video-on-demand, individuals can select pre-recorded content from a central information bank, allowing a movie or film clip to be broadcasted immediately when requested. Common examples of VOD include Netflix Watch Instantly movies, Hulu television shows, training videos, and news broadcasts.

A ProxySG stores frequently requested data and distributes it upon client requests. Because the ProxySG is closer to the client than the origin server, the data is served locally, which saves bandwidth and increases quality of service by reducing pauses or buffering during playback. Because of its proximity to the end user, the ProxySG provides higher quality streams (also dependent on the client connection rate) than the origin server.

Note that VOD content can be paused, rewound, and played back.

Serving Content: Multicast Streaming

Multicast transmission is analogous to a radio frequency on which any device can listen. Any device that supports multicast can transmit on the multicast channel. One copy of the data is sent to a group address. Devices in the group listen for traffic at the group address and join the stream if clients in the routing tree are requesting the stream. Only the group participants receive the traffic at the address associated with the group. Broadcasts differ from multicast because broadcast traffic is sent to the entire network.

For multicast transmission to occur, the network devices through which the content is to be sent must support multicast. In particular:

- ❑ Content creators must explicitly set up their streaming servers to support multicast.

For example, for Windows Media, content creators can set up multicast-enabled stations, stations that are not multicast-enabled, or both. For RealNetworks, the configuration of the server includes specifying whether the server supports multicast and, if so, which clients (subnets) can use multicast.

- ❑ Routers on the path must support multicast.
- ❑ Clients must request a multicast transmission. Media players that are set for multicast transmission simply join the multicast channel to receive the streaming data, sometimes without establishing an explicit one-to-one connection to the device sending the transmission.

Benefits of Multicast

The benefits of using multicast for streaming media include the following:

- ❑ It alleviates network congestion.
- ❑ For live streaming events that have a large audience, multicast significantly reduces network traffic compared to the traffic that would result from transmitting the same live event over unicast. If unicast transport is used, the same content must be sent across the network multiple times or it must be broadcast to all devices on the network.

- ❑ It scales well as the number of participants expand.
- ❑ It is well suited for efficient transmission over satellite links.

A company might, for example want to reserve WAN connections for business-critical traffic, such as stock trades, but it needs a way to deliver corporate broadcasts. The company could efficiently transmit corporate broadcasts over satellite by using multicast transmission and reserve the WAN for business-critical traffic.
- ❑ It enables network planners to proactively manage network growth and control cost because deploying multicast is more cost-effective than alternatives for increasing LAN and WAN capabilities.

Limitations of Multicast

The limitations of multicast include the following:

- ❑ Multicast support is not yet widely available on the Internet. Therefore, using multicast to deliver content is limited to intranet-style deployments.
- ❑ Not all networking equipment supports multicasting. In addition, not all network administrators enable the multicast functionality on their networking equipment.
- ❑ Switches do not understand multicast. When a multicast stream reaches a switch, the switch sends the multicast stream to all of its ports. A switch treats a multicast address as an Ethernet broadcast.

About Serving Multicast Content

The ProxySG takes a multicast stream from the origin server and delivers it as a unicast stream. This avoids the main disadvantage of multicasting—that all of the routers on the network must be multicast-enabled to accept a multicast stream. Unicast-to-multicast, multicast-to-multicast, and broadcast alias-(scheduled live from stored content)-to-multicast are also supported.

For Windows Media multicast, a Windows Media Station file (.NSC) is downloaded through HTTP to acquire the control information required to set up content delivery.

For Real Media, multicasting maintains a TCP control (accounting) channel between the client and media server. The multicast data stream is broadcast using UDP from the ProxySG to streaming clients, who join the multicast.

Limiting Bandwidth

The following sections describe how to configure the ProxySG to limit global and protocol-specific media bandwidth.

To manage streaming media bandwidth, you configure the ProxySG to restrict the total number of bits per second the appliance receives from the origin media servers and delivers to clients. The configuration options are flexible to allow you to configure streaming bandwidth limits for the ProxySG, as well as for the streaming protocol proxies (Windows Media, Real Media, and QuickTime).

Note: Bandwidth claimed by HTTP, non-streaming protocols, and network infrastructure is not constrained by this limit. Transient bursts that occur on the network can exceed the hard limits established by the bandwidth limit options.

After it has been configured, the ProxySG limits streaming access to the specified threshold. If a client tries to make a request after a limit has been reached, the client receives an error message.

Note: If a maximum bandwidth limitation has been specified for the ProxySG, the following condition can occur. If a Real Media client, followed by a Windows Media client, requests streams through the same ProxySG and total bandwidth exceeds the maximum allowance, the Real Media client enters the rebuffering state. The Windows Media client continues to stream.

Consider the following features when planning to limit streaming media bandwidth:

- ❑ ProxySG to server (all protocols)—The total kilobits per second allowed between the appliance and any origin content server or upstream proxy for all streaming protocols. Setting this option to 0 effectively prevents the ProxySG from initiating any connections to the media server. The ProxySG supports partial caching in that no bandwidth is consumed if portions of the media content are stored in the ProxySG.

Limiting ProxySG bandwidth restricts the following streaming media-related functions:

- Live streaming, where the proxy requests from the server, the sum of all unique bit rates requested by the clients
 - The ability to fetch new data for an object that is partially cached
 - Reception of multicast streams
- ❑ Client to ProxySG (all protocols)—The total kilobits per second allowed between streaming clients and the ProxySG. Setting this option to 0 effectively prevents any streaming clients from initiating connections through the ProxySG.

Limiting client bandwidth restricts the following streaming media-related functions:

- MBR support; when lower bit-rate selection by the client could have allowed the client to stream, the client is denied when the bandwidth limit is exceeded
- Limits the transmission of multicast streams

- ❑ **Client connections**—The total number of clients that can connect concurrently. When this limit is reached, clients attempting to connect receive an error message and are not allowed to connect until other clients disconnect. Setting this variable to 0 effectively prevents any streaming media clients from connecting.

Selecting a Method to Limit Streaming Bandwidth

The ProxySG offers two methods for controlling streaming bandwidth. The way that each method controls bandwidth differs—read the information below to decide which method best suits your deployment requirements.

Limiting streaming bandwidth using the streaming features (described in this chapter) works as follows: if a new stream comes in that pushes above the specified bandwidth limit, that new stream is denied. This method allows existing streams to continue to get the same level of quality they currently receive.

The alternate way of limiting streaming bandwidth is with the bandwidth management feature. With this technique, all streaming traffic for which you have configured a bandwidth limit shares that limit. If a new stream comes in that pushes above the specified bandwidth limit, that stream is allowed, and the amount of bandwidth available for existing streams is reduced. This causes streaming players to drop to a lower bandwidth version of the stream. If a lower bandwidth version of the stream is not available, players that are not receiving enough bandwidth can behave in an unpredictable fashion. In other words, if the amount of bandwidth is insufficient to service all of the streams, some or all of the media players experience a reduction in stream quality. For details, see ["Bandwidth Management"](#) on page 597.

Because of the degradation in quality of service, for most circumstances, Blue Coat recommends that you use the streaming features to control streaming bandwidth rather than the bandwidth management features. Do *not* use both methods at the same time.

Caching Behavior: Proxy Specific

This section describes the type of content the ProxySG caches for each supported proxy.

Flash

The ProxySG caches pre-recorded audio and video content delivered over Real Time Messaging Protocol (RTMP) or RTMP traffic tunneled over HTTP (RTMPT). Flash media files have `.flv`, `.f4v` extensions.

MS Smooth

The ProxySG caches on-demand Smooth Streaming video content delivered over HTTP. Silverlight is the typical player used for Smooth Streaming and is available as a plug-in for web browsers running under Microsoft Windows and Mac OS X.

Windows Media

The ProxySG caches Windows Media-encoded video and audio files. The standard extensions for these file types are: `.wmv`, `.wma`, and `.asf`.

Real Media

The ProxySG caches Real Media-encoded files, such as RealVideo and RealAudio. The standard extensions for these file types are: `.ra`, `.rm`, and `.rmvb`. Other content served from a Real Media server through RTSP is also supported, but it is not cached. This content is served in *pass-through* mode only. (Pass-through mode offers application, layer-7 proxy functionality, but does not support acceleration features—caching, pre-population, splitting, and multi-casting.)

QuickTime

The ProxySG does not cache QuickTime content (`.mov` files). All QuickTime content is served in pass-through mode only.

Adobe HDS

The ProxySG caches on-demand and live video content delivered over HTTP.

Adobe HLS

The ProxySG caches on-demand and live video content delivered over HTTP.

Caching Behavior: Video-on-Demand

The ProxySG supports the caching of files for VOD streaming. First, the client connects to the ProxySG, which in turn connects to the origin server and pulls the content, storing it locally. Subsequent requests of this same content are served from the ProxySG. This provides bandwidth savings, as every hit to the ProxySG means less network traffic. Blue Coat also supports partial caching of streams.

Note: On-demand files must be unicast.

Splitting Behavior: Live Broadcast

The ProxySG supports splitting of live content, but behavior varies depending upon the media type.

For live streams, the ProxySG can split streams for clients that request the same stream. First, the client connects to the ProxySG, which then connects to the origin server and requests the live stream. Subsequent requests of the same content from different clients are split from the appliance.

Two streams are considered identical by the ProxySG if they share the following characteristics:

- ❑ The stream is a live or broadcast stream.
- ❑ The URL of the stream requested by client is identical.

- ❑ MMS (Microsoft Media Services), MMSU (MMS UDP), and MMST (MMS TCP) are considered to be identical.
- ❑ RTMP and RTMPT are considered to be identical.

Splitting of live unicast streams provides bandwidth savings, since subsequent requests do not increase network traffic.

Multiple Bit Rate Support

Content authors normally encode streaming media content into different bit rates to meet the needs of the different speeds of Internet access—modem, ISDN, DSL, and LAN. In contrast, the delivery bit rate is the actual speed at which the content is delivered to the client. For example, a stream encoded for playback at 56Kbps must be delivered to clients at a bit rate of 56Kbps or higher. A client with enough bandwidth might ask the streaming server to send the 56Kbps encoded stream at 220Kbps; the data is buffered locally and played back at 56Kbps. The playback experience of 56Kbps stream delivered at 220Kbps would be better at 220Kbps than at 56Kbps. The reason is that more time is available for the client to request packets to be retransmitted if packets are dropped.

The ProxySG supports multiple bit rate (MBR), which is the capability of a single stream to deliver multiple bit rates to clients requesting content from caches from within varying levels of network conditions (such as different connecting bandwidths and varying levels of competing traffic). MBR allows the ProxySG and the client to negotiate the optimal stream quality for the available bandwidth even when the network conditions are bad. MBR increases client-side streaming quality, especially when the requested content is not cached.

The ProxySG caches only the requested bit rate. For example, a media client that requests a 50Kbps stream receives that stream, and the ProxySG caches only the 50Kbps bit rate content, no other rate.

Flash has a similar functionality called *dynamic streaming*. Like MBR, dynamic streaming allows clients to switch to a bitrate suitable for current network conditions.

Note: The Flash proxy does not cache videos that the OCS delivers by dynamic streaming.

Bit Rate Thinning

Thinning support is closely related to MBR, but thinning allows for data rate optimizations even for single data-rate media files. If the media client detects that there is network congestion, it requests a subset of the single data rate stream. For example, depending on how congested the network is, the client requests only the *key video frames* or audio-only instead of the complete video stream.

Pre-Populating Content

Note: This feature applies to Windows Media and Real Media only.

The ProxySG supports pre-population of streaming files from both HTTP (Web) servers and origin content servers (that is, streaming servers). Downloading streaming files from HTTP servers reduces the time required to pre-populate the file.

Note: Smooth Streaming and QuickTime content cannot be pre-populated.

Pre-population can be accomplished through streaming from the media server. The required download time is equivalent to the file length; for example, a two-hour movie requires two hours to download. Now, if the media file is hosted on an HTTP server, the download time occurs at normal transfer speeds of an HTTP object, and is independent of the play length of the media file.

Note: Content must be hosted on an HTTP server in addition to the media server.

Using the `content distribute` CLI command, content is downloaded from the HTTP server and renamed with a given URL argument. A client requesting the content perceives that the file originated from a media server. If the file on the origin media server experiences changes (such as naming convention), SGOS bypasses the cached mirrored version and fetches the updated version.

Example:

```
content distribute rtsp://wm_server/bar.wmv from http://web_server/
bar.wmv
```

Note: In the example above, `rtsp://wm_server/bar.wmv` should also be accessible as a streaming object on a streaming server.

About Fast Streaming (Windows Media)

Note: This feature applies to Windows Media only.

Windows Media Server version 9 and higher contains a feature called Fast Streaming that allows clients to provide streams with extremely low buffering time.

SGOS supports the following functionality for both cached and uncached content:

- ❑ Fast Start—Delivers an instant playback experience by eliminating buffering time. The first few seconds of data are sent using the maximum available bandwidth so that playback can begin as soon as possible.

- ❑ **Fast Cache**—Streams content to clients faster than the data rate that is specified by the stream format. For example, fast caching allows the server to transmit a 128-kilobits-per-second (Kbps) stream at 500 Kbps. The Windows Media client buffers the streaming content before it is rendered at the specified rate — 128 Kbps for this stream.

In the case of MBR VOD content, fast-caching content to the local cache of the Windows Media client impacts playback quality. To maintain smooth streaming of MBR VOD content, you might need to disable the fast-caching ability of the Windows Media client. By default, fast-caching is enabled on the ProxySG. You can use the VPM or CPL to configure policy for disabling fast caching, thereby preventing the Windows Media clients from fast-caching content to the local cache. For the VPM and CPL properties, see the *Visual Policy Manager Reference* and the *Content Policy Language Reference*.

Fast Recovery and Fast Reconnect are currently not supported on the ProxySG.

About QoS Support

The ProxySG supports Quality of Service (QoS), which allows you to create policy to examine the Type of Service fields in IP headers and perform an action based on that information. For streaming protocols, managing the QoS assists with managing bandwidth classes.

For detailed information about managing QoS, see the Advanced Policy chapter in *Visual Policy Manager Reference*.

IPv6 Support

All streaming proxies include IPv6 support, and the ProxySG can act as a transitional device between IPv4 and IPv6 networks for Flash, Smooth Streaming over HTTP, Windows Media (RTSP, HTTP), Real Media, and QuickTime. Streaming proxies support IPv6 in the following ways:

- ❑ **Flash:** RTMP-based protocols (such as RTMP, RTMPT) support IPv6 for making upstream connections to the origin content server (OCS) as well as can accept IPv6 client connections.
- ❑ **MS Smooth, Adobe HDS, and Apple HLS:** Protocols streaming over HTTP support IPv6 for making upstream connections to the OCS, and can accept IPv6 client connections.
- ❑ **Windows Media:**
 - RTSP and HTTP protocols support IPv6 for making upstream connections to the OCS, and can accept IPv6 client connections.
 - For multicast-station, the RTSP protocol can be used when retrieving content from an IPv6 OCS and sending multicast to IPv4 clients.
 - ASX rewrite is IPv6 capable, but only for the HTTP protocol.

- ❑ **Real Media and QuickTime:** RTSP and HTTP protocols support IPv6 for making upstream connections to the OCS, and can accept IPv6 client connections.

Note that Windows Media over MMS does not support IPv6.

About Streaming Media Authentication

The following sections discuss authentication between streaming media clients and ProxySGs and between ProxySGs and origin content servers (streaming servers).

Flash Proxy Authentication

The RTMP protocol does not include support for challenge/response authentication. For RTMP traffic tunneled over HTTP (RTMPT), proxy authentication is done by the HTTP proxy, without involvement of the Flash proxy; this is true regardless of whether the handoff to Flash proxy is enabled or disabled.

If an authenticate (<realm>) policy involves challenging the user, those RTMP connections will be denied access.

MS Smooth Proxy Authentication

Because Smooth Streaming uses HTTP as its transport protocol, all proxy authentication options supported for HTTP are supported for Smooth Streaming. These proxy authentication options will exhibit the same behavior regardless of whether the HTTP handoff for Smooth Streaming is enabled.

Windows Media Server-Side Authentication

Windows Media server authentication for HTTP and MMS supports the following authentication types:

- ❑ HTTP—BASIC Authentication and Membership Service Account
- ❑ HTTP—BASIC Authentication and Microsoft Windows Integrated Windows Authentication (IWA) Account Database
- ❑ IWA Authentication and IWA Account Database

The ProxySG supports the caching and live-splitting of server-authenticated data. It has partial caching functionality so that multiple security challenges are not issued to Windows Media Player when it accesses different portions of the same media file.

The first time Windows Media content is accessed on the streaming server, the ProxySG caches the content along with the authentication type that was enabled on the origin server at the time the client sent a request for the content. The cached authentication type remains until the appliance learns that the server has changed

the enabled authentication type, either through cache coherency (checking to be sure the cached contents reflect the original source) or until the ProxySG connects to the origin server (to verify access credentials).

Windows Media Proxy Authentication

If you configure proxy authentication on the ProxySG, Windows Media clients are authenticated based on the policy settings. The ProxySG evaluates the request from the client and verifies the accessibility against the set policies. Windows Media Player then prompts the client for the proper password. If the client password is accepted, the Windows Media server might also require the client to provide a password for authentication. If a previously accepted client attempts to access the same Windows Media content again, the ProxySG verifies the user credentials using its own credential cache. If successful, the client request is forwarded to the Windows Media server for authentication.

Windows Media Player Authentication Interactivities

Consider the following proxy authentication interactivities with Windows Media Player (except when specified, these do not apply to HTTP streaming):

- ❑ If the proxy authentication type is configured as BASIC and the server authentication type is configured as IWA, the default is denial of service.
- ❑ If proxy authentication is configured as IWA and the server authentication is configured as BASIC, the proxy authentication type defaults to BASIC.
- ❑ The ProxySG does not support authentication based on `url_path` or `url_path_regex` conditions when using `mms` as the `url_scheme`.
- ❑ Transparent style HTTP proxy authentication fails to work with Windows Media Players when the credential cache lifetime is set to 0 (independent of whether server-side authentication is involved).
- ❑ If proxy authentication is configured, a request for a stream through HTTP prompts the user to enter access credentials twice: once for the proxy authentication and once for the media server authentication.
- ❑ Additional scenarios involving HTTP streaming exist that do not work when the TTL is set to zero (0), even though only proxy authentication (with no server authentication) is involved. The ProxySG returning a 401-style proxy authentication challenge to Windows Media Player 6.0 does not work because the Player cannot resolve inconsistencies between the authentication response code and the server type returned from the ProxySG. This results in an infinite loop of requests and challenges. Example scenarios include transparent authentication—resulting from either a transparent request from a player or a hard-coded service specified in the ProxySG—and request of cache-local (ASX-rewritten or unicast alias) URLs.

Windows Media Server Authentication Type (MMS)

Note: This section applies to Windows Media MMS and requires the CLI.

Configure the ProxySG to recognize the type of authentication the origin content server is using: BASIC or NTLM/Kerberos.

To configure the media server authentication type for WM-MMS:

At the (config) prompt, enter the following command:

```
SGOS#(config) streaming windows-media server-auth-type {basic | ntlm}
```

Real Media Proxy Authentication

If you configure proxy authentication on the ProxySG, Real Media clients are authenticated based on the policy settings. The ProxySG evaluates the request from the client and verifies the accessibility against the set policies. Next, RealPlayer prompts the client for the proper password. If the client password is accepted, the Real Media server can also require the client to provide a password for authentication. If a previously accepted client attempts to access the same Real Media content again, the ProxySG verifies the user credentials using its own credential cache. If successful, the client request is forwarded to the Real Media server for authentication.

Real Media Player Authentication Limitation

Using RealPlayer 8.0 in transparent mode with both proxy and Real Media server authentication configured to BASIC, RealPlayer 8.0 always sends the same proxy credentials to the media server. This is regardless of whether a user enters in credentials for the media server. Therefore, the user is never authenticated and the content is not served.

QuickTime Proxy Authentication

BASIC is the only proxy authentication mode supported for QuickTime clients. If an IWA challenge is issued, the mode automatically downgrades to BASIC.

Adobe HDS Authentication

Because Adobe HDS uses HTTP as its transport protocol, all proxy authentication options supported for HTTP are supported for Adobe HDS. These proxy authentication options will exhibit the same behavior regardless of whether the HTTP handoff for Adobe HDS is enabled.

Apple HLS Authentication

Because Apple HLS uses HTTP as its transport protocol, all proxy authentication options supported for HTTP are supported for Apple HLS. These proxy authentication options will exhibit the same behavior regardless of whether the HTTP handoff for Apple HLS is enabled.

Section B: Configuring Streaming Media

This section describes how to configure the various ProxySG streaming options. It contains the following topics:


- ❑ "Configuring the HTTP Streaming Proxy" on page 550
- ❑ "Configuring Streaming Services to Intercept Traffic" on page 551
- ❑ "Configuring the Windows Media, Real Media, and QuickTime Proxies" on page 554
- ❑ "Limiting Bandwidth" on page 556
- ❑ "Configuring the ProxySG Multicast Network" on page 558
- ❑ "Forwarding Client Logs" on page 558
- ❑ "Reference: Access Log Fields" on page 559
- ❑ "Reference: CPL Triggers, Properties, and Actions for Streaming Proxies" on page 561
- ❑ "Viewing Streaming History Statistics" on page 562

Configuring the HTTP Streaming Proxy

To optimize streaming over HTTP, you need to intercept the HTTP services used for Smooth Streaming, Adobe HDS, and Apple HLS traffic, and configure the corresponding proxy to accept hand off from the HTTP proxy. For additional information, see "About HTTP-Based Streaming" on page 533.

To intercept the HTTP services:

1. From the Management Console, select **Configuration > Services > Proxy Services**.



2. Change the applicable HTTP services to Intercept:
 - a. In the **Standard** service group, locate the applicable HTTP service: **Explicit HTTP** or **External HTTP**.

Note: The **Internal HTTP** service is set to use the TCP Tunnel proxy—not the HTTP proxy. If you have an internal MS Smooth server whose traffic you want to optimize, Symantec recommends creating a new service that intercepts the traffic from that streaming server.

- b. Select **Intercept** for each set of ports defined for the service.
3. Click **Apply**.

To configure the HTTP streaming proxy:

1. From the Management Console, select **Configuration > Proxy Settings > Streaming Proxies**.
2. Select the **HTTP** tab.



3. **Enable Microsoft Smooth Streaming, Adobe HTTP Dynamic Streaming, and Apple HTTP Live Streaming handoff:** Enabled by default. When an HTTP Streaming client requests a stream through the ProxySG, the HTTP proxy service passes control to the appropriate proxy, so that HTTP streaming will be supported through the HTTP proxy port. Disable one of these options only if you do not want to optimize traffic for that protocol.
4. Click **Apply**.

Configuring Streaming Services to Intercept Traffic


By default (upon upgrade and on new systems), the ProxySG has streaming services configured on ports 1755 (MMS) and 554 (RTSP). In addition to port 1935 (RTMP), ports 8080/ 80 (Explicit HTTP) can also be used for Flash applications. The services are configured to listen to all IP addresses, but are set to **Bypass** mode.

To configure streaming services to intercept Flash media-based traffic, see ["Using the Flash Streaming Proxy"](#) on page 584.

The following procedure describes how to change the service to **Intercept** mode.

To configure the MMS/RTSP proxy services attributes:

1. From the Management Console, select **Configuration > Services > Proxy Services**.



2. Change the streaming services to Intercept:
 - a. Scroll through the list of services and select the **Standard** service group; select the **MMS** and **RTSP** groups.
 - b. From the **MMS All ->All:1755** row drop-down list, select **Intercept**.
 - c. From the **RTSP All ->All:554** row drop-down list, select **Intercept**.

3. Click **Apply**.

Now that the streaming listeners are configured, you can configure the streaming proxies. Proceed to:


- ❑ ["Configuring the Windows Media, Real Media, and QuickTime Proxies"](#) on page 554 to configure the proxy options that determine how to process streaming traffic.
- ❑ (Optional) ["Adding a New Streaming Service"](#) (below) to add new streaming services that bypass specific network segments or listen on ports other than the defaults.

Adding a New Streaming Service


The ProxySG allows you to add new streaming services. Consider the following scenario: you want the ProxySG to exclude (bypass) an IP address/subnet from intercepting streaming traffic because that network segment is undergoing routine maintenance.

To add a new streaming service:

1. From the Management Console, select **Configuration > Services > Proxy Services**.



2. Scroll the list of services and select the **Standard** service group.
3. Click **New Service**. The New Service dialog displays with the default settings.



4. Configure the service options:
 - a. Name the service. In this example, the service is named **ExcludeStreaming** because the network admin wants to prevent the ProxySG from intercepting streaming traffic from a specific IP address.
 - b. From the **Service Group** drop-down list, select **Standard**—the service group to which streaming traffic belongs.
 - c. From the **Proxy** drop-down list, select **MMS**, **RTSP**, or **RTMP**.

Note: To bypass traffic from multiple streaming protocols, create another service for the streaming protocol not selected in this step.


- d. Click **New**. The New Listener dialog displays.
- e. This example selects the **Destination host or subnet** option and specifies a sample IP address.
- f. This example accepts the default value of **554**, the default port for the RTSP protocol. If the ProxySG is intercepting streaming traffic on a different port, you must specify the port number here.
- g. This example selects **Bypass** as the option; the ProxySG will not intercept streaming traffic.
- h. Click **OK** in each dialog to close them.

Configuring the Windows Media, Real Media, and QuickTime Proxies

This section describes how to configure the Windows Media, Real Media, and QuickTime proxies. The Windows Media and Real Media proxy options are identical except for one extra option for Real Media. QuickTime has only one option (**Enable HTTP Handoff**).

To configure Windows Media, Real Media, and QuickTime streaming proxies:

1. From the Management Console, select **Configuration > Proxy Settings > Streaming Proxies**.
2. Select the tab for the proxy you want to configure: **Windows Media**, **Real Media**, **QuickTime**.



3. **Enable HTTP handoff:** Enabled by default. When a Windows Media, Real Media, or QuickTime client requests a stream from the ProxySG over port 80, which in common deployments is the only port that allows traffic through a firewall, the HTTP module passes control to the streaming module so HTTP streaming can be supported through the HTTP proxy port. Disable this option only if you do not want HTTP streams to be cached or split.
4. **Forward client-generated logs to origin media server:** Enabled by default. The ProxySG logs information, such as client IP address, the date, and the time, to the origin server for Windows Media and Real Media content. See ["Forwarding Client Logs"](#) on page 558 for more information about log forwarding.
5. **Enable multicast (Real Media proxy only):** The ProxySG receives a unicast stream from the origin RealServer and serves it as a multicast broadcast. This allows the ProxySG to take a one-to-one stream and split it into a one-to-many stream, saving bandwidth and reducing the server load. It also produces a higher quality broadcast.

Multicasting maintains a TCP control (accounting) channel between the client and RealServer. The multicast data stream is broadcast using UDP from the ProxySG to RealPlayers that join the multicast. The ProxySG support for Real Media uses UDP port 554 (RTSP) for multicasting. This port number can be changed to any valid UDP port number.
6. Specify how often the ProxySG checks cached streaming content for freshness.
 - **Never check freshness:** Although this is the default setting, Blue Coat recommends selecting one of the other freshness options.
 - **Check freshness every *value* hours:** The ProxySG checks content freshness every *n.nn* hours.

Note: A value of 0 requires the streaming content to always be checked for freshness.

- **Check freshness every access:** Every time cached content is requested, it is checked for freshness.
7. Configure bandwidth limit options:
 - To limit the bandwidth for client connections to the ProxySG, select **Client bandwidth limit (kbits/sec)**. In the **Kbits/sec** field, enter the maximum number of kilobits per second that the ProxySG allows for all streaming client connections.
 - To limit the bandwidth for connections from the ProxySG to origin content servers, select **Gateway bandwidth limit (kbits/sec)**. In the **kbits/sec** field, enter the maximum number of kilobits per second that the ProxySG allows for all streaming connections to origin media servers.
 8. To limit the bandwidth for connections from the ProxySG to the OCS, select **Client Connections Limit**. In the **clients** field, enter the total number of clients that can connect concurrently.
 9. Click **Apply**.

Note: For multicast, additional configuration is required. See ["Configuring the ProxySG Multicast Network"](#) on page 558.

See Also

- ["Configuring Streaming Services to Intercept Traffic"](#)
- ["Limiting Bandwidth"](#)
- ["Managing Multicast Streaming for Windows Media"](#)
- ["Managing Simulated Live Content \(Windows Media\)"](#)
- ["Windows Media Player Interactivity Notes"](#)

Limiting Bandwidth

This section describes how to limit bandwidth from the clients to the ProxySG and from the ProxySG to origin content servers.


Configuring Bandwidth Limits—Global

This section describes how to limit bandwidth use of Windows Media, Real Media, and QuickTime streaming protocols through the ProxySG.

Note: This global setting does not control Flash or Smooth Streaming traffic.

To specify the global bandwidth limit for streaming protocols:

1. Select **Configuration > Proxy Settings > Streaming Proxies > General**.



2. To limit the client connection bandwidth:
 - a. In the **Bandwidth** field, select **Client bandwidth limit (kbits/sec)**. In the **kbits/sec** field, enter the maximum number of kilobits per second that the ProxySG allows for all streaming client connections.

Note: This option is not based on individual clients.

- b. In the **Bandwidth** pane, select **Gateway bandwidth limit (kbits/sec)**. In the **kbits/sec** field, enter the maximum number of kilobits per second that the ProxySG allows for all streaming connections to origin media servers.

3. Click **Apply**.

See Also

- ❑ ["Configuring Streaming Services to Intercept Traffic" on page 551](#)
- ❑ ["Configuring the Windows Media, Real Media, and QuickTime Proxies" on page 554](#)
- ❑ ["Configuring the ProxySG Multicast Network" on page 558](#)
- ❑ ["Viewing Streaming History Statistics" on page 562](#)

Configuring Bandwidth Limitation—Fast Start (WM)

Note: This section applies to Windows Media only and requires the CLI.

Upon connection to the ProxySG, Windows Media clients do not consume more bandwidth (in kilobits per second) than the defined value.

To specify the maximum starting bandwidth:

At the (config) prompt, enter the following command:

```
SGOS#(config) streaming windows-media max-fast-bandwidth kbps
```

Limiting Bandwidth for Smooth Streaming

The global bandwidth limits for streaming protocols do not apply to Smooth Streaming because it is essentially just HTTP traffic. However, you can write policy to limit bandwidth of Smooth Streaming clients:

```
<proxy>
streaming.client=ms_smooth limit_bandwidth.client_outbound(bw_class)
```

Configuring the ProxySG Multicast Network

This section describes how to configure the ProxySG multicast service. Additional steps are required to configure the ProxySG to serve multicast broadcasts to streaming clients (Windows Media and Real Media); those procedures are provided in subsequent sections.

To configure the multicast service:

1. Select **Configuration > Proxy Settings > Streaming Proxies > General**.

The screenshot shows the 'General' tab of the ProxySG configuration interface. The 'Multicast' section is expanded, showing the following fields:

- Maximum hops:** 16 (Callout 2a points to this field)
- IP range:** 224.2.128.0 to 224.2.255.255 (Callout 2b points to the start IP)
- Port range:** 32768 to 65535 (Callout 2c points to the start port)

2. Configure multicast options:
 - a. In the **Maximum hops** field, enter a time-to-live (TTL) value.
 - b. In the **IP range** fields, enter the range of IP addresses that are available for multicast.
 - c. In the **Port range** fields, enter the range of ports available for multicast.
3. Click **Apply**.
4. Enable multicast:
 - Real Media: See [Step 5](#) on page 555.
 - Windows Media: See ["Managing Multicast Streaming for Windows Media"](#) on page 565.

Forwarding Client Logs

The ProxySG can log information about Windows Media and Real Media streaming sessions between the client and the ProxySG appliance and can also forward these client-generated logs to the origin media server. Additionally, for Windows Media RTSP only, ProxySG appliance also supports forwarding values for certain fields to the server, when windows-media streaming proxy has log forwarding enabled and logging compatibility disabled.

Note: For Real Media, the log is only forwarded before a streaming session is halted; QuickTime log forwarding is not supported.

The following fields are included in the client log record:

- ❑ `cs-uri-stem`: URI stem of the client request.
- ❑ `s-cpu-util`: CPU utilization of the ProxySG.
- ❑ `s-totalclients`: Clients connected to the ProxySG (but not necessarily receiving streams).
- ❑ `s-pkts-sent`: Number of packets the ProxySG sent to the client, during the playspurt.
- ❑ `s-proxied`: Set to 1 for proxied sessions.
- ❑ `s-session-id`: A unique ID of the streaming session between the client and the ProxySG.
- ❑ `sc-bytes`: Number of bytes the ProxySG sent to the client, during the playspurt.

To enable/disable log forwarding:

Use the Management Console (see "[Configuring the Windows Media, Real Media, and QuickTime Proxies](#)" on page 554) or use the following CLI command at the (config) prompt:

```
SGOS#(config) streaming windows-media log-forwarding {enable |
disable}
```

To enable/disable RTSP log compatibility:

At the (config) prompt, enter the following command:

```
SGOS#(config) streaming windows-media log-compatibility {enable |
disable}
```

Reference: Access Log Fields

Two streaming log formats are available: *streaming* and *bcreporterstreaming_v1*. To see which format is being used for streaming, select **Configuration > Access Logging > Logs > Logs**; the format is listed next to the name. To change the format used for the log, go to the **General Settings** tab.

Legacy Streaming Log Format

The legacy streaming log format contains the following fields:

```
c-ip date time c-dns cs-uri-scheme cs-host cs-uri-port cs-uri-path cs-
uri-query c-starttime x-duration c-rate c-status c-playerid c-
playerversion c-playerlanguage cs(User-Agent) cs(Referer) c-hostexe c-
hostexever c-os c-osversion c-cpu filelength filesize avgbandwidth
protocol transport audiocodec videocodec channelURL sc-bytes c-bytes
```

```
s-pkts-sent c-pkts-received c-pkts-lost-client c-pkts-lost-net c-pkts-lost-cont-net c-resendreqs c-pkts-recovered-ECC c-pkts-recovered-resent c-buffercount c-totalbuffertime c-quality s-ip s-dns s-totalclients s-cpu-util x-cache-user s-session-id x-cache-info x-client-address s-action
```

The streaming-specific access log fields are described below, in alphabetical order.

- ❑ **audiocodec:** Audio codec used in the stream.
- ❑ **channelURL:** URL to the .nsc file.
- ❑ **c-buffercount:** Number of times the client buffered while playing the stream.
- ❑ **c-bytes:** An MMS-only value of the total number of bytes delivered to the client.
- ❑ **c-playerid:** Globally unique identifier (GUID) of the player.
- ❑ **c-playerlanguage:** Client language-country code.
- ❑ **c-playerversion:** Version number of the player.
- ❑ **c-rate:** Mode of Windows Media Player when the last command event was sent.
- ❑ **c-starttime:** Timestamp (in seconds) of the stream when an entry is generated in the log file.
- ❑ **c-totalbuffertime:** Time (in seconds) the client used to buffer the stream.
- ❑ **protocol:** Protocol used to access the stream: mms, http, asfm, rtsp, rtmp, rtmpt, rtmpe, rtmpte.
- ❑ **s-session-id:** Session ID for the streaming session.
- ❑ **s-totalclients:** Clients connected to the server (but not necessarily receiving streams).
- ❑ **transport:** Transport protocol used (UDP, TCP, multicast, and so on).
- ❑ **videocodec:** Video codec used to encode the stream.
- ❑ **x-cache-info:** Values: UNKNOWN, DEMAND_PASSTHRU, DEMAND_MISS, DEMAND_HIT, LIVE_PASSTHRU, LIVE_SPLIT.
- ❑ **x-duration:** Length of time a client played content prior to a client event (FF, REW, Pause, Stop, or jump to marker).
- ❑ **x-wm-c-dns:** Hostname of the client determined from the Windows Media protocol.
- ❑ **x-wm-c-ip:** The client IP address determined from the Windows Media protocol.
- ❑ **x-cs-streaming-client:** Type of streaming client in use (windows_media, real_media, quicktime, flash, ms_smooth).
- ❑ **x-rs-streaming-content:** Type of streaming content served (windows_media, real_media, quicktime, flash). Note that ms_smooth (Smooth Streaming over HTTP) is not a possible value for this field.

- ❑ `x-streaming-bitrate`: The reported client-side bitrate for the stream.

Reporter Streaming Log Format

The `bcreporterstreaming_v1` log format contains the following fields:

```
date time time-taken c-ip sc-status s-action sc-bytes rs-bytes cs-
method cs-uri-scheme cs-host cs-uri-port cs-uri-path cs-uri-query cs-
username cs-auth-group cs(Referer) cs(User-Agent) c-starttime
filelength filesize avgbandwidth x-rs-streaming-content x-streaming-
rtmp-app-name x-streaming-rtmp-stream-name x-streaming-rtmp-swf-url x-
streaming-rtmp-page-url s-ip s-dns s-session-id x-cache-info
```

The streaming-specific access log fields are described below.

- ❑ `x-rs-streaming-content`: Type of streaming content served (`windows_media`, `real_media`, `quicktime`, `flash`). Note that `ms_smooth` (Smooth Streaming over HTTP) is not a possible value for this field.
- ❑ `x-streaming-rtmp-app-name`: The application parameter in an RTMP "connect" command. In VOD, it usually corresponds to a directory on the OCS file system.
- ❑ `x-streaming-rtmp-stream-name`: Name of the stream requested by the Flash client. In VOD, it often corresponds to a filename in the OCS file system.
- ❑ `x-streaming-rtmp-swf-url`: URL of the Flash client SWF file (if sent in the RTMP connect request)
- ❑ `x-streaming-rtmp-page-url`: URL of the web page in which the Flash client SWF file is embedded (if sent in the RTMP connect request)

When encrypted streaming protocols are tunneled (either because of policy or because the protocol version is unknown) some of the information can not be ascertained. The following fields are not available for RTMPE or RTMPTE video sites when the encrypted connection is tunneled:

```
c-starttime
filelength
filesize
avgbandwidth
x-streaming-rtmp-app-name
x-streaming-rtmp-stream-name
x-streaming-rtmp-swf-url
x-streaming-rtmp-page-url
```

Reference: CPL Triggers, Properties, and Actions for Streaming Proxies

The following Blue Coat CPL is supported in all streaming proxies. For Flash-specific CPL triggers and properties, see ["Reference: CPL Triggers and Properties for Flash"](#) on page 589.

Triggers

```
streaming.client=
streaming.content= (not applicable to MS Smooth proxy)
```

Properties and Actions

`streaming.fast_cache()` (Windows Media proxy only)

`streaming.transport()` (not applicable to MS Smooth proxy)


Viewing Streaming History Statistics

The **Streaming History** tabs display bar graphs that illustrate the number of active client connections over the last hour (60 minutes), day (24 hours), and month (30 days) for a specific streaming proxy (Windows Media, Real Media, QuickTime, and Flash). These statistics are not available through the CLI. The **Current Streaming Data** and **Total Streaming Data** tabs display real-time values for current connections and live traffic activity on the ProxySG. Current and total streaming data statistics are available through the CLI.

Note: The MS Smooth (Smooth Streaming) proxy does not currently collect data to be displayed in the streaming history panel.

To view client statistics:

1. Select **Statistics > Protocol Details > Streaming History**.
2. Select the client type for which you want to view statistics under the **Protocol** drop down menu: **Windows Media**, **RealMedia**, **QuickTime**, and **Flash**.
3. Select the **Duration**: from the drop-down menu.
Choose from **Last Hour**, **Last Day**, **Last Month**, and **All Periods**.



4. (Optional) To set the graph scale to a different value, select a value from the **Graph scale should** drop-down list.


Viewing Current and Total Streaming Data Statistics

The Management Console **Current Streaming Data** tab and the **Total Streaming Data** tab show real-time values for Windows Media, Real Media, QuickTime, and Flash activity on the ProxySG. These statistics can also viewed using the CLI.

Viewing Current Streaming Data Statistics

To view current streaming data statistics:

1. Select **Statistics > Protocol Details > Streaming History > Current Streaming Data**.




2. Select a streaming protocol (**Windows Media, Real Media, QuickTime, Flash**) from the **Protocol** drop-down list.
3. Select a traffic connection type (**Live Traffic, On-Demand Traffic, or Passthru Traffic**) from the drop-down list.

Viewing Total Streaming Data Statistics

To view total streaming data statistics:

1. Select **Statistics > Streaming History > Total Streaming Data**.



2. Select a streaming protocol (**Windows Media, Real Media, QuickTime, Flash**) from the **Protocol** drop-down list.

3. Select a traffic connection type (**Live Traffic**, **On-Demand**, or **Passthru Traffic**) from the drop-down list.

To clear streaming statistics:

To zero-out the streaming statistics, enter the following command at the CLI prompt:

```
SGOS# clear-statistics {quicktime | real-media | windows-media}
```

Note: The `clear-statistics` command cannot be used to clear Flash statistics.

Section C: Additional Windows Media Configuration Tasks

This section provides Windows Media configuration tasks that aren't available through the Management Console, but can be executed through the CLI.

This section contains the following topics:

- ❑ "Managing Multicast Streaming for Windows Media" on page 565
- ❑ "Managing Simulated Live Content (Windows Media)" on page 569
- ❑ "ASX Rewriting (Windows Media)" on page 571

Managing Multicast Streaming for Windows Media

This section describes multicast station and `.nsc` files, and explains how to configure the ProxySG to send multicast broadcasts to Windows Media clients.

See the following sections:

- ❑ "About Multicast Stations"
- ❑ "Creating a Multicast Station"
- ❑ "Monitoring the Multicast Station"
- ❑ "Multicast to Unicast Live Conversion at the ProxySG"
- ❑ "Managing Multicast Streaming for Windows Media"

About Multicast Stations

A *multicast station* is a defined location from where Windows Media Player retrieves live streams. This defined location allows Advanced Streaming Format (`.asf`) streams to be delivered to many clients using only the bandwidth of a single stream. Without a multicast station, streams must be delivered to clients through unicast.

A multicast station contains all of the information needed to deliver `.asf` content to a Windows Media Player or to another ProxySG, including:

- ❑ IP address
- ❑ Port
- ❑ Stream format
- ❑ TTL value (time-to-live, expressed hops)

The information is stored in an `.nsc` file, which Window Media Player must be able to access to locate the IP address.

If Windows Media Player fails to find proper streaming packets on the network for multicast, the player can roll over to a unicast URL. Reasons for this include lack of a multicast-enabled router on the network or if the player is outside the multicast station's TTL. If the player fails to receive streaming data packets, it uses the unicast URL specified in the `.nsc` file. All `.nsc` files contain a unicast URL to allow rollover.

Unicast to Multicast

Unicast to multicast streaming requires converting a unicast stream on the server-side connection to a multicast station on the ProxySG. The unicast stream must contain live content before the multicast station works properly. If the unicast stream is a video-on-demand file, the multicast station is created but is not able to send packets to the network. For video-on-demand files, use the `broadcast-alias` command. A *broadcast alias* defines a playlist, and specifies a starting time, date, and the number of times the content is repeated.

Multicast to Multicast

Use the `multicast-alias` command to get the source stream for the multicast station.

Creating a Multicast Station

To create a multicast station, you perform the following steps:

- ❑ Define a name for the multicast station.
- ❑ Define the source of the multicast stream.
- ❑ (Optional) Change the port range to be used.
- ❑ (Optional) Change the IP address range of the multicast stream.
- ❑ (Optional) Change the Time-to-Live (TTL) value. TTL is a counter within an ICMP packet. As a packet goes through each router, the router decrements this TTL value by 1. If the packet traverses enough routers for the value to reach 0, routers will no longer forward this packet.

Note: For MMS protocol only, you can use an alias—`multicast-alias`, `unicast-alias`, or `broadcast-alias`—as a source stream for a multicast station. WM-RTSP and WM-HTTP do not support aliases.

Syntax

```
multicast-station name {alias | url} [address | port | ttl]
```

where

- *name* specifies the name of the multicast station, such as `station1`.
- {*alias* | *url*} defines the source of the multicast stream. The source can be a URL or it can be a multicast alias, a unicast alias, or simulated live. (The source commands must be set up before the functionality is enabled within the multicast station.)
- [*address* | *port* | *ttl*] are optional commands that you can use to override the default ranges of these values. (Defaults and permissible values are discussed below.)

Example 1: Create a Multicast Station

This example:

- ❑ Creates a multicast station, named `station1`, on ProxySG `10.25.36.47`.
- ❑ Defines the source as `rtsp://10.25.36.47/tenchi`
- ❑ Accepts the address, port, and TTL default values.

```
SGOS#(config) streaming windows-media multicast-station station1
rtsp://10.25.36.47/tenchi.
```

To delete multicast `station1`:

```
SGOS#(config) streaming no multicast-station station1
```

Example 2: Create a Broadcast Alias and Direct a Multicast Station to Use it as the Source

This example:

- ❑ To allow unicast clients to connect through multicast, creates a broadcast alias named `array1`; defines the source as `mms://10.25.36.48/tenchi2`.
- ❑ Instructs the multicast station from Example 1, `station1`, to use the broadcast alias, `array1`, as the source.

```
SGOS#(config) streaming windows-media broadcast-alias array1 mms://
10.25.36.48/tenchi2 0 today noon
SGOS#(config) streaming windows-media multicast-station station1
array1
```

Changing Address, Port, and TTL Values

Specific commands allow you to change the address range, the port range, and the default TTL value. To leave the defaults as they are for most multicast stations and change it only for specified station definitions, use the `multicast-station` command.

The `multicast-station` command randomly creates an IP address and port from the specified ranges.

- ❑ Address-range: the default ranges from `224.2.128.0` to `224.2.255.255`; the permissible range is between `224.0.0.2` and `239.255.255.255`.
- ❑ Port-range: the default ranges from `32768` to `65535`; the permissible range is between `1` and `65535`.
- ❑ TTL value: the default value is `5` hops; the permissible range is from `1` to `255`.

Syntax, with Defaults Set

```
multicast address-range <224.2.128.0>-<224.2.255.255>
multicast port-range <32768>-<65535>
multicast ttl <5>
```

Getting the .nsc File

The `.nsc` file is created from the multicast station definition and saved through the browser as a text file encoded in a Microsoft proprietary format.

Without an `.nsc` file, the multicast station definition does not work.

To create an `.nsc` file from the newly created `station1`, open the file by navigating through the browser to the multicast station's location (where it was created) and save the file as `station1.nsc`.

The file location, based on the streaming configuration above:

```
http://10.25.36.47/MMS/nsc/station1.nsc
```

Save the file as `station1.nsc`.

Note: You can also enter the URL in Windows Media Player to start the stream.

The newly created file is not editable; the settings come from the streaming configuration file. In that file, you have already defined the following pertinent information:

- ❑ The address, which includes TTL, IP address, IP port, Unicast URL, and the NSC URL. All created `.nsc` files contain a unicast URL for rollover in case Windows Media Player cannot find the streaming packets.
- ❑ The description, which references the RTSP URL that you defined.
- ❑ The format, which contains important Advanced Streaming Format (ASF) header information. All streams delivered by the multicast station definition have their ASF headers defined here.

Monitoring the Multicast Station

You can determine the multicast station definitions by viewing the streaming Windows Media configuration.

To view the multicast station setup:

```
SGOS#(config) show streaming windows config
; Windows Media Configuration
license: 1XXXXXXX-7XXXXXXX-7XXXXX
logging: enable
logging enable
http-handoff: enable
live-retransmit: enable
transparent-port (1755): enable
explicit proxy: 0
refresh-interval: no refresh interval (Never check freshness)
max connections: no max-connections (Allow maximum
connections)
max-bandwidth: no max-bandwidth (Allow maximum bandwidth)
max-gateway-bandwidth: no max-gateway-bandwidth (Allow maximum
bandwidth)
multicast address: 224.2.128.0 - 224.2.255.255
multicast port: 32768 - 65535
multicast TTL: 5
asx-rewrite: No rules
```

```
multicast-alias:          No rules
unicast-alias:            No rules
broadcast-alias:          No rules
multicast-station:        station1 rtsp://10.25.36.47/tenchi
224.2.207.0 40465 5 (playing)
```

Note: *Playing* at the end of the multicast station definition indicates that the station is currently sending packets onto the network. The IP address and port ranges have been randomly assigned from the default ranges allowed.

To determine the current client connections and current ProxySG connections, use the `show streaming windows-media statistics` command.

To view the multicast station statistics:

```
SGOS#(config) show streaming windows stat
;Windows Media Statistics
Current client connections:
  by transport: 0 UDP, 0 TCP, 0 HTTP, 1 multicast
  by type: 1 live, 0 on-demand
Current gateway connections:
  by transport: 0 UDP, 1 TCP, 0 HTTP, 0 multicast
  by type: 1 live, 0 on-demand
```

Multicast to Unicast Live Conversion at the ProxySG

The ProxySG supports converting multicast streams from an origin content server to unicast streams. The stream at the ProxySG is given the appropriate unicast headers to allow the appliance to direct one copy of the content to each user on the network.

Multicast streaming only uses UDP protocol and does not know about the control channel, which transfers essential file information. The `.nsc` file (a file created off-line that contains this essential information) is retrieved at the beginning of a multicast session from an HTTP server. The `multicast-alias` command specifies an alias to the URL to receive this `.nsc` file.

The converted unicast stream can use any of the protocols supported by Windows Media, including HTTP streaming.

When a client requests the alias content, the ProxySG uses the URL specified in the `multicast-alias` command to fetch the `.nsc` file from the HTTP server. The `.nsc` file contains all of the multicast-related information, such as addresses and `.asf` file header information that is normally exchanged through the control connection for unicast-delivered content.

Note: For Windows Media streaming clients, additional multicast information is provided in ["Managing Multicast Streaming for Windows Media"](#) on page 565.

Managing Simulated Live Content (Windows Media)

This section describes simulated live content and how to configure the ProxySG to manage and serve simulated live content.

About Simulated Live Content

The simulated live content feature defines playback of one or more video-on-demand files as a scheduled live event, which begins at a specified time. The content can be looped multiple times, or scheduled to start at multiple start times throughout the day. If used in conjunction with the `multicast-alias` command, the live content is multicast; otherwise, live content is accessible as live-splitting sources. The feature does *not* require the content to be cached.

When you have set a starting date and time for the simulated live content, the broadcast of the content starts when at least one client requests the file. Clients connecting during the scheduled playback time of the simulated live content receive cached content for playback. Clients requesting the simulated live content before the scheduled time are put into wait mode. Clients requesting the content after all of the contents have played receive an error message. Video-on-demand content does not need to be on the ProxySG before the scheduled start time, but pre-populating the content on the provides better streaming quality.

The ProxySG computes the starting playtime of the broadcast stream based on the time difference between the client request time and the simulated live starting time.

Before configuring simulated live, consider the following:

- ❑ The simulated live content name must be unique. Aliases are not case sensitive.
- ❑ The name cannot be used for both a unicast and a multicast alias name.
- ❑ After simulated live content is referenced by one or more multicast stations, the simulated live content cannot be deleted until all multicast stations referencing the simulated live content are first deleted.

The multicast station appears as another client of simulated live content, just like a Windows Media Player.

Note: This note applies to HTTP only. If a client opens Windows Media Player and requests an alias before the starting time specified in the broadcast-alias option, the HTTP connection closes after a short time period. When the specified time arrives, the player fails to reconnect to the stream and remains in waiting mode.

Creating a Broadcast Alias for Simulated Live Content

Syntax

```
streaming windows-media broadcast-alias alias url loops date time
```

where:

- *alias* is the name of the simulated live content.
- *url* is the URL for the video-on-demand stream. Up to 128 URLs can be specified for simulated live content.

- *loops* is the number of times you want the content to be played back. Set to 0 (zero) to allow the content to be viewed an indefinite number of times.
- *date* is the simulated live content starting date. Valid date strings are in the format *yyyy-mm-dd* or *today*. You can specify up to seven start dates by using the comma as a separator (no spaces).
- *time* is the simulated live content starting time. Valid time strings are in the format *hh:mm* (on a 24-hour clock) or one of the following strings:
 - midnight, noon
 - 1am, 2am, ...
 - 1pm, 2pm, ...

Specify up to 24 different start times within a single date by using the comma as a separator (no spaces).

Example 1

This example creates a playlist for simulated live content. The order of playback is dependent on the order you enter the URLs. You can add up to 128 URLs.

```
SGOS#(config) streaming windows-media broadcast-alias alias url
```

Example 2

This example demonstrates the following:

- ❑ creates a simulated live file called *bca*.
- ❑ plays back *rtsp://ocs.bca.com/bca1.asf* and *rtsp://ocs.bca.com/bca2.asf*.
- ❑ configures the ProxySG to play back the content twice.
- ❑ sets a starting date and time of today at 4 p.m., 6 p.m., and 8 p.m.

```
SGOS#(config) streaming windows-media broadcast-alias bca rtsp://  
ocs.bca.com/bca1.asf 2 today 4pm,6pm,8pm  
SGOS#(config) streaming windows-media broadcast-alias bca rtsp://  
ocs.bca.com/bca2.asf
```

To delete simulated live content:

```
SGOS#(config) streaming windows-media no broadcast-alias alias
```

ASX Rewriting (Windows Media)

This section describes ASX rewriting and applies to Windows Media only.

An ASX file is an active streaming redirector file that points to a Windows Media audio or video presentation. It is a metafile that provides information about Active Streaming Format (ASF) media files.

See the following topics:

- ❑ ["About ASX Rewrite"](#)
- ❑ ["Windows Media Player Interactivity Notes"](#)
- ❑ ["Configuring the Windows Media, Real Media, and QuickTime Proxies"](#)

About ASX Rewrite

If your environment does not use a Layer 4 switch or the Cisco Web Cache Control Protocol (WCCP), the ProxySG can operate as a proxy for Windows Media Player clients by rewriting the Windows Media ASX file (which contains entries with URL links to the actual location of the streaming content) to point to the ProxySG rather than the Windows Media server.

The metadata files can have `.asx`, `.wvx`, or `.wax` extensions, but are commonly referred to as ASX files. The ASX file references the actual media files (with `.asf`, `.wmv`, and `.wma` extensions). An ASX file can refer to other `.asx` files, although this is not a recommended practice. If the file does not have one of the metafile extensions and the Web server that is serving the metadata file does not set the correct MIME type, it is not processed by the Windows Media module. Also, the `.asx` file with the appropriate syntax must be located on an HTTP (not a Windows Media) server.

The ASX rewrite module is triggered by either the appropriate file extension or the returned MIME type from the server (`x-video-asf`).

Note: If an `.asx` file syntax does not follow the standard `<ASX>` tag-based syntax, the ASX rewrite module is not triggered.

For the ProxySG to operate as a proxy for Windows Media Player requires the following:

- ❑ The client is explicitly proxied for HTTP content to the ProxySG that rewrites the `.asx` metafile.
- ❑ The streaming media ProxySG is configurable.

Note: Windows Media Player automatically tries to roll over to different protocols according to its Windows Media property settings before trying the rollover URLs in the `.asx` metafile.

With the `asx-rewrite` command, you can implement redirection of the streaming media to a ProxySG by specifying the rewrite protocol, the rewrite IP address, and the rewrite port.

The protocol specified in the ASX rewrite rule is the protocol the client uses to reach the ProxySG. You can use forwarding and policy to change the default protocol specified in the original `.asx` file that connects to the origin media server.

When creating ASX rewrite rules, you need to determine the number priority. It is likely you will create multiple ASX rewrite rules that affect the `.asx` file; for example, rule 100 could redirect the IP address from 10.25.36.01 to 10.25.36.47, while rule 300 could redirect the IP address from 10.25.36.01 to 10.25.36.58. In this case, you are saying that the original IP address is redirected to the IP address in rule 100. If that IP address is not available, the ProxySG looks for another rule matching the incoming IP address.

Notes and Interactivities

Before creating rules, consider the following.

- ❑ Each rule you create must be checked for a match; therefore, performance might be affected if you create many rules.
- ❑ Low numbers have a higher priority than high numbers.

Note: You must use the CLI to create rule.

- ❑ ASX rewrite rules configured for multiple ProxySGs configured in an HTTP proxy-chaining configuration can produce unexpected URL entries in access logs for the *downstream* ProxySG (the ProxySG to which the client proxies). The combination of proxy-chained ProxySGs in the HTTP path coupled with ASX rewrite rules configured for multiple ProxySGs in the chain can create a rewritten URL requested by the client in the example form of:

```
protocol1://downstream_SecApp/redirect?protocol2://<upstream_
SecApp>/redirect?protocol3://origin_host/origin_path
```

In this scenario, the URL used by the downstream ProxySG for caching and access logging can be different than what is expected. Specifically, the downstream ProxySG creates an access log entry with `protocol2://upstream_SecApp/redirect` as the requested URL. Content is also cached using this truncated URL. Blue Coat recommends that the ASX rewrite rule be configured for only the downstream ProxySG, along with a proxy route rule that can forward the Windows Media streaming requests from the downstream to upstream ProxySGs.

Syntax for the `asx-rewrite` Command

```
asx-rewrite rule # in-addr cache-proto cache-addr [cache-port]
```

where:

- `in-addr`—Specifies the hostname or IP address delivering the content
- `cache-proto`—Specifies the rewrite protocol on the ProxySG. Acceptable values for the rewrite protocol are:
 - `mmsu` specifies Microsoft Media Services UDP
 - `mmst` specifies Microsoft Media Services TCP
 - `http` specifies HTTP

- `mms` specifies either MMS-UDP or MMS-TCP
- `*` specifies the same protocol as in the `.asx` file

If the `.asx` file is referred from within another `.asx` file (not a recommended practice), use a `*` for the `cache-proto` value. The `*` designates that the protocol specified in the original URL be used. As a conservative, alternative approach, you could use HTTP for the `cache-proto` value.

- `cache-addr`—Specifies the rewrite address on the ProxySG.
- `cache-port`—Specifies the port on the ProxySG. This value is optional.

To set up the `.asx` rewrite rules:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) streaming windows-media asx-rewrite number in-addr  
cache-proto cache-addr cache-port
```

Note: To delete a specific rule, enter `streaming windows-media no asx-rewrite number`.

To ensure that an ASX rewrite rule is immediately recognized, clear the local browser cache.

Example

This example:

- ❑ Sets the priority rule to 200.
- ❑ Sets the protocol to be whatever protocol was originally specified in the URL and directs the data stream to the appropriate default port.
- ❑ Provides the rewrite IP address of 10.9.44.53, the ProxySG.

```
SGOS#(config) streaming windows-media asx-rewrite 200 * * 10.9.44.53
```

Note: ASX files must be fetched from HTTP servers. If you are not sure of the network topology or the content being served on the network, use the asterisks to assure the protocol set is that specified in the URL.

ASX Rewrite Incompatibility With Server-side IWA Authentication

Server-side authentication (MMS only, not HTTP) is supported if the origin media server authentication type is BASIC or No Auth. However, if you know that a Windows Media server is configured for IWA authentication, the following procedure allows you to designate any virtual IP addresses to the IWA authentication type. If you know that all of the activity through the ProxySG requires IWA authentication, you can use the IP address of the appliance.

To designate an IP address to an authentication type:

1. If necessary, create a virtual IP address that is used to contact the Windows Media server.
2. At the (config) prompt, enter the following command:
`SGOS#(config) streaming windows-media server-auth-type ntlm ip_address`
3. Configure the ASX rewrite rule to use the IP address.
 - a. To remove the authentication type designation:
`SGOS#(config) streaming windows-media no server-auth-type
ip_address`
 - b. To return the authentication type to BASIC:
`SGOS#(config) streaming windows-media server-auth-type basic
ip_address`

Section D: Configuring Windows Media Player


This section describes how to configure Windows Media Player to communicate through the ProxySG.

To apply the ProxySG Windows Media streaming services, Windows Media Player must be installed and configured to use explicit proxy. For a transparent deployment, no WMP configuration is necessary.

Note: The following procedure example uses Windows Media Player 11. Installation and setup varies with different versions of Windows Media Player.

To configure Windows Media Player:

1. Start Windows Media Player.
2. Select **Tools > Options**.



3. Navigate to protocol configuration:
 - a. Select **Network**.
 - b. Select **HTTP**.
 - c. Click **Configure**. The Configure Protocol dialog displays.

4. Configure the proxy settings:
 - a. Select **Use the following proxy server**.
 - b. Enter the ProxySG IP address and the port number used for the explicit proxy (the default HTTP port is 80). These settings must match the settings configured in the ProxySG. If you change the ProxySG explicit proxy configuration, you must also reconfigure Windows Media Player.
5. Click **OK** in both dialogs. Result: Windows Media Player now proxies through the ProxySG and content is susceptible to streaming configurations and access policies.

Windows Media Player Interactivity Notes

This section describes Windows Media Player inter activities that might affect performance.

Striding

When you use Windows Media Player, consider the following interactivities in regard to using fast forward and reverse (referred to as *striding*):

- ❑ If you request a cached file and repeatedly attempt play and fast forward, the file freezes.
- ❑ If you attempt a fast reverse of a cached file that is just about to play, you receive an error message, depending on whether you have a proxy:
 - Without a proxy: A device attached to the system is not functioning.
 - With a proxy: The request is invalid in the current state.
- ❑ If Windows Media Player is in pause mode for more than ten minutes and you press fast reverse or fast forward, an error message displays: `The network connection has failed.`

Other Notes

- ❑ Applies to WMP v9: If a `url_host_rewrite` rule is configured to rewrite a host name that is a domain name instead of an IP address, a request through the MMS protocol fails and the host is not rewritten. As the connect message sent by the player at the initial connection does not contain the host name, a rewrite cannot occur. HTTP requests are not affected by this limitation.
- ❑ If explicit proxy is configured and the access policy on the ProxySG is set to `deny`, a requested stream using HTTP from Windows Media Player 9 serves the stream directly from the origin server even after the request is denied. The player sends a request to the OCS and plays the stream from there.

Blue Coat recommends the following policy:

```
<proxy>
    streaming.content=yes deny
-or-
<proxy>
    streaming.content=windows_media deny
```

The above rules force the HTTP module to hand off HTTP requests to the MMS module. MMS returns the error properly to the player, and does not go directly to the origin server to try to serve the content.

- ❑ If you request an uncached file using the HTTP protocol, the file is likely to stop playing if the authentication type is set to BASIC or NTLM/Kerberos and you initiate rapid seeks before the buffering begins for a previous seek. Windows Media Player, however, displays that the file is still playing.
- ❑ If a stream is scheduled to be accessible at a future time (using a simulated live rule), and the stream is requested before that time, Windows Media Player enters a waiting stage. This is normal. However, if HTTP is used as the protocol, after a minute or two Windows Media Player closes the HTTP connection, but remains in the waiting stage, even when the stream is broadcasting.

Notes:

For authentication-specific notes, see "[Windows Media Server-Side Authentication](#)" on page 547 and "[Windows Media Proxy Authentication](#)" on page 548.

Section E: Configuring RealPlayer


This section describes how to configure Real Player to communicate through the ProxySG.

To use the ProxySG Real Media streaming services with an explicit proxy configuration, the client machine must have RealPlayer installed and configured to use RTSP streams. If you use transparent proxy, no changes need to be made to RealPlayer.

Note: This procedure features RealPlayer, version 10.5. Installation and setup menus vary with different versions of RealPlayer. Refer to the RealPlayer documentation to configure earlier versions of RealPlayer.

To configure RealPlayer:

1. Start RealPlayer.
2. Select **Tools > Preferences**.




3. Navigate to proxy settings:
 - a. Select **Connection > Proxy**.
 - b. Click **Change Settings**. The Streaming Proxy Settings dialog appears.
4. Configure options:
 - a. In the **PNA and RTSP proxies:** field, select **Use proxies**.
 - b. Enter the ProxySG IP address and the port number used for the explicit proxy (the default RTSP port is 544). These settings must match the settings configured in the ProxySG. If you change the ProxySG explicit proxy configuration, you must also reconfigure RealPlayer. If using transparent proxy, RTSP port 554 is set by default and cannot be changed.

Note: For **HTTP Proxy**, if you have an HTTP proxy already configured in your browser, select **Use system Internet Connection proxy settings**.

- c. Optional: For **HTTP Proxy**, if you have an HTTP proxy already configured in your browser, select **Use system Internet Connection proxy settings**.
- d. Optional: In the **Do not use proxy for:** section, you can enter specific hosts and bypass the ProxySG.


Note: This can also be accomplished with policy, the method Blue Coat recommends.

- e. Click **OK** to close the Streaming Proxy Settings dialog.



5. Configure RealPlayer transport settings:
 - a. Select **Connection > Network Transports**.
 - b. Click **RTSP Settings**. The RTSP Transport Settings dialog displays.
6. If required, deselect options, based on your network configuration. For example, if your firewall does not accept UDP, you can deselect **Attempt to use UDP for all content**, but leave the TCP option enabled. Blue Coat recommends using the default settings.
7. Click **OK**.

To allow the creation of access log entries, RealPlayer must be instructed to communicate with the RealServer.



8. Perform the following:

- a. **Select View > Preferences > Internet/Privacy.**
- b. In the **Privacy** field, select **Send connection-quality data to RealServers**; click **OK**.


Result: RealPlayer now proxies through the ProxySG and content is susceptible to streaming configurations and access policies.

Note: For authentication-specific issues, see "[Real Media Proxy Authentication](#)" on page 549.

Section F: Configuring QuickTime Player

This section describes how to configure QuickTime player for explicit proxy to the ProxySG.

1. Start QuickTime player.
2. Select **Edit > Preferences > QuickTime Preferences**.



3. Configure the protocol settings:
 - a. Click **Advanced**.
 - b. Select **RTSP Proxy Server**;
 - c. Enter the IP address of the ProxySG.
 - d. Enter the port number (554 is the default).

These settings must match the settings configured in the ProxySG. If you change the ProxySG explicit proxy settings, set similar settings in QuickTime.

4. Close **OK**.

Result: QuickTime now proxies—in pass-through mode—through the ProxySG.

Note: For authentication-specific issues, see "[QuickTime Proxy Authentication](#)" on page 549.

Section G: Using the Flash Streaming Proxy

This section describes how to use the Flash streaming proxy.

- ❑ ["Configuring the Flash Streaming Proxy"](#) on page 584
- ❑ ["Additional Information"](#) on page 587
- ❑ ["Reference: CPL Triggers and Properties for Flash"](#) on page 589
- ❑ ["Reference: VPM Reference Information."](#) on page 591

Configuring the Flash Streaming Proxy

Note: The Flash streaming proxy requires a valid Flash license.

Perform these tasks to configure the Flash proxy so that it splits live streams and caches video-on-demand.

Task #	Task	Reference
1	Configure the client browsers to use the ProxySG appliance as an explicit proxy. Required for explicit deployments only.	"Configuring Client Browsers for Explicit Proxy" on page 584.
2	Intercept the RTMP service on transparent deployments. or Intercept the Explicit-HTTP service on explicit deployments.	"Intercepting the RTMP Service (Transparent Deployment)" on page 585 or "Intercepting the Explicit HTTP Service (Explicit Deployment)" on page 585
3	Enable HTTP handoff so that RTMP tunneled over HTTP is also intercepted.	"Enabling HTTP Handoff for the Flash Proxy" on page 586
4	Verify optimization of Flash traffic.	"Verifying Optimization of Flash Traffic" on page 586

Configuring Client Browsers for Explicit Proxy

To set up the Web browser manually, you must include the following information in the Internet Explorer browser configuration:

- ❑ The fully-qualified hostname or IP address of the ProxySG appliance. You cannot use a hostname only.
- ❑ The port on which the appliance will listen for traffic. The default port is 8080.


Note: You cannot configure Firefox browsers because Flash uses Windows settings.

1. In Internet Explorer, select **Tools > Internet Options**.

2. Select the **Connections** tab.
3. If you are using a LAN, click **LAN Settings**. If you are using a Dial-up or Virtual Private Network connection, click **Add** to set up the connection wizard.
4. Make sure the **Automatically detect proxy settings** and **Use a proxy automatic configuration script** options are *not* checked.
5. Select **Use a proxy server** for your LAN.
6. Select **Advanced**. The Proxy Settings dialog displays.
7. For HTTP, enter the IP address of the ProxySG appliance, and add the port number; 8080 is the default.
8. Select **Use the same proxy server for all protocols**.
9. Click **OK** and exit out of all open dialogs.

Intercepting the RTMP Service (Transparent Deployment)

To optimize Flash traffic in a transparent deployment, you need to have an RTMP proxy service configured to listen on port 1935 (the typical RTMP port), and this service must be set to intercept. This service also controls RTMPE traffic.



Most likely, you will already have an RTMP service; if not, you should create it. Then set the service to Intercept:


1. In the Management Console, select **Configuration > Services > Proxy Services**.
2. Locate the RTMP service in the Standard group.
3. Select **Intercept**.
4. Click **Apply**.

Intercepting the Explicit HTTP Service (Explicit Deployment)

To optimize Flash traffic in an explicit deployment, you should have an Explicit HTTP proxy service configured to listen on ports 8080 and 80, and this service must be set to intercept. This service controls plain and encrypted Flash connections tunneled over HTTP.

Most likely, you will already have an Explicit HTTP service; if not, you should create it. Then set the service to Intercept:

1. In the Management Console, select **Configuration > Services > Proxy Services**.
2. Locate the Explicit HTTP service in the Standard group.




3. Select **Intercept** for Explicit:8080 and Explicit:80.
4. Click **Apply**.

Enabling HTTP Handoff for the Flash Proxy

If Flash clients are unable to connect over raw RTMP due to firewall restrictions, the players are sometimes configured to tunnel RTMP over HTTP (RTMPT). In order to intercept and cache content that uses the RTMPT protocol, you need to enable the HTTP handoff for the Flash proxy. Starting in SGOS 6.2.x, the HTTP handoff is enabled by default on *new* installations, but you may need to enable the setting on *upgraded* systems, if you haven't already done so.

1. In the Management Console, select **Configuration > Proxy Settings > Streaming Proxies**.
2. Select the **Flash** tab.




3. Select the **Enable HTTP handoff** check box and click **Apply**.

Verifying Optimization of Flash Traffic

When a live stream is being split, or a stream in a VOD connection is being cached or is played from the cache, the Active Sessions report shows an in-color Object Caching (OC) icon . The following steps show you how to verify caching of a pre-recorded video.

1. Using a Flash client, play a pre-recorded video (one that you have not played previously).
2. While the video is playing, go to the Management Console and select **Statistics > Sessions > Active Sessions**.
3. For **Filter**, select **Proxy** and choose **Flash**.
4. Click **Show** to display a list of connections.
5. Locate the connection. It is listed with Flash in the **Protocol** column and an in-color Object Caching  icon in the **OC** column since the connection is being cached. The **Savings** column indicates little to no bandwidth savings since this is the first time the video was played.
6. Play the same video again.
7. Display the active Flash proxy sessions. Because the video was served from the cache, there is significant bandwidth savings shown in the **Savings** column.



Client IP	Server	A	S	FW	I	Duration	Client Bytes	Server Bytes	Savings	C	BC	OC	P	BM	Service Name	Protocol
10.99.99.11:57113	10.99.99.151...					15.8 min	193,336	361,852	0%						Exploit HTTP	Flash
10.99.99.11:57158	10.99.99.151...					13.4 min	113,175	443,568	0%						Exploit HTTP	Flash
10.99.99.11:57159	10.99.99.151...					13.3 min	19,124,444	443,483	97.88%						Exploit HTTP	Flash
10.99.99.11:57180	10.99.99.151...					12.5 min	2,230,705	378,820	83.02%						Exploit HTTP	Flash
10.99.99.11:57183	10.99.99.151...					12.3 min	6,753,424	587,220	91.32%						Exploit HTTP	Flash
10.99.99.11:57255	10.99.99.151...					8.8 min	2,228,609	340,620	84.71%						Exploit HTTP	Flash
10.99.99.11:57255	10.99.99.151...					8.1 min	143,166	200,190	0%						Exploit HTTP	Flash
10.99.99.11:57289	10.99.99.151...					7.2 min	6,752,609	487,732	92.77%						Exploit HTTP	Flash
10.99.99.11:57298	10.99.99.151...					6.9 min	2,228,620	322,612	85.53%						Exploit HTTP	Flash
10.99.99.11:57341	10.99.99.151...					5 min	2,230,281	282,682	87.33%						Exploit HTTP	Flash
10.99.99.11:57353	10.99.99.151...					4.7 min	83,952	121,382	0%						Exploit HTTP	Flash
10.99.99.11:57376	10.99.99.151...					3.3 min	6,630,208	223,665	96.63%						Exploit HTTP	Flash
10.99.99.11:57381	10.99.99.151...					3.2 min	2,230,550	234,676	89.48%						Exploit HTTP	Flash
10.99.99.11:57403	10.99.99.151...					2 min	2,596,339	235,103	90.94%						Exploit HTTP	Flash

Encrypted Flash connections will show one of the following three messages in the **Detail** column:

- ❑ **Encrypted**—The encrypted connection was decrypted, optimized, and re-encrypted.
- ❑ **Encrypted, tunneled by policy**—The encrypted connection was not decrypted or optimized because a policy dictated that the connection should be tunneled. The policy property that controls whether encrypted Flash connections are tunneled is `streaming.rtmp.tunnel_encrypted()`.
- ❑ **Encrypted, tunneled as unknown protocol version**—The encrypted connection could not be decrypted or optimized because the RTMPE protocol version was not recognized.


Additional Information

See the following sections for additional information related to the Flash proxy:

- ❑ ["When VOD Content Gets Cached"](#) on page 588
- ❑ ["Proxy Chaining"](#) on page 588
- ❑ ["CDN Interoperability Support"](#) on page 589

When VOD Content Gets Cached

- ❑ The Flash proxy caches fully-played and partially-played portions of VOD content. If a video is played from the beginning to the end, the file is fully cached. If a video is stopped in the middle of play, only the played portion is cached. The next time a user requests the same video, the cached portion will be served from the cache and the remainder of the video will be played from the OCS (and added to the cache).
- ❑ With the default settings:
 - If the playing video is not already cached, it will be cached as it is played from the OCS.
 - If the playing video has already been cached, it will be played from the cache.
 - If the playing video is stored in the cache but the cache is out of date from what is on the OCS, it will not play from the cache or be written to cache.
- ❑ The Flash proxy caches content that is connected to the beginning of the video (User 1 and User 2 below). If a playspurt isn't attached to the beginning of the video, the content cannot be cached (User 3.) In order for content to be appended to the cache, the client must begin playing the video somewhere from within the cached region; then, when the uncached content is played from the OCS, it will be added to the cache (User 2).



- ❑ Encrypted and plain content are stored separately in the object cache.

Proxy Chaining

Proxy chaining (hierarchy of proxies) supports the use of multiple ProxySG appliances between the server and client. This hierarchy of proxy servers (set by the administrator using policy gestures) allows further maximizing of bandwidth usage optimization achieved by features such as live splitting. If forwarding is set up in an organized manner, the overhead involved in splitting and transmitting live streams gets pushed to the end of the proxy chain (the one closest to the end users), which avoids sending any piece of content across any given WAN link more than once.

To enable proxy chaining, you must create forwarding hosts using the MC or CLI and set the proxy hierarchy using the following policy gestures:

- ❑ Traffic can be forwarded to the next ProxySG appliance by using the following policy gesture:

```
forward(fwd_host)
```

where the *fwd_host* must be type `proxy` and have a defined `http` port.

- ❑ Traffic can be forwarded to the *server* by using the following policy gesture:

```
forward(fwd_host)
```

where the *fwd_host* must be type `server` and have a defined `rtmp` port.

Use the following CLI command to create forwarding hosts:

```
#(config forwarding)create host ?
<host-alias> <host-name> [http[=<port>]] [https[=<port>]]
[ftp[=<port>]]
[mms[=<port>]] [rtmp[=<port>]] [rtsp[=<port>]] [tcp=<port>]
[telnet[=<port>]]
[ssl-verify-server[=(yes|no)]] [group=<group-name>] [server|proxy]
```

CDN Interoperability Support

To maximize performance within forward proxy deployments using CDNs, the Flash proxy supports interoperability with the following features:

- ❑ **SWF verification:** Support for SWF verification by the Flash Media server.
- ❑ **FCSubscribe/FCUnsubscribe, onFCSubscribe/onFCUnsubscribe:** Interoperability support for these messages used by some CDNs for live streams. (not applicable to VOD caching)
- ❑ **Use of Ident services:** Support to ensure bandwidth optimization from splitting is preserved even when using Ident service.
- ❑ **Token-based authentication:** Support for relaying authorization information between clients and servers.

Note: A server connection is maintained on behalf of each client connection due to CDN interoperability reasons.

Reference: CPL Triggers and Properties for Flash

Flash streaming proxy supports policy enforcement based on RTMP traffic. The tables below lists the CPL commands for generating Flash streaming policy. For more information about the CPL, refer to the *Content Policy Language Guide*.

The Flash-related CPL triggers are listed below:

Table 24–5 CPL Triggers

CPL Trigger	Supported Values
client.protocol	rtmp, rtmpt, rtmpe, rtmpte
request.header.User-Agent	<string>
streaming.client	flash
streaming.rtmp.app_name	<string>

Table 24–5 CPL Triggers

CPL Trigger	Supported Values
streaming.rtmp.method	open, connect, play
streaming.rtmp.page_url	<URL>
streaming.rtmp.stream_name	<string>
streaming.rtmp.swf_url	<URL>
url	<URL>
live	yes, no
streaming.content	flash

The CPL properties related to Flash are listed below:

Table 24–6 CPL Properties

CPL Property	Comments
access_server	This property is ignored for Flash VOD caching because the Flash proxy <i>always</i> checks the OCS for every playspurt.
allow, deny, force_deny	
always_verify	This property is ignored for Flash VOD caching because the Flash proxy will always verify object requests with the OCS. Therefore, even in fully-cached videos, you will see some server bytes statistics.
bypass_cache	Traffic is enforced on a per-stream basis and not the entire application. If this property is set to yes, the video is played directly from the server even if the content is cached. If set to no (the default), cached portions of the video play from the cache and uncached portions play from the OCS.
cache	This setting is overridden by <code>bypass_cache (yes)</code> . If this property is set to yes (the default), VOD content is cached. If set to no and the file is fully cached, the video is played from the cache. If set to no and the file is not cached or is partially cached, the video is played in pass-through mode.
delete_on_abandonment	This property is ignored for Flash VOD caching since it's not applicable.
force_cache	This property is ignored for Flash VOD caching since the RTMP protocol does not have any headers that indicate cacheability.
forward	Forwarding to http hosts of type proxy and http and rtmp hosts of type server allowed.
forward.fail_open	

Table 24-6 CPL Properties

CPL Property	Comments
max_bitrate	Not supported for Flash.
reflect_ip	
streaming.rtmp.tunnel_encrypted	Determines whether encrypted Flash traffic (RTMPE and RTMPTE) is tunneled or accelerated.
streaming.transport	streaming.transport(http) can be used to coerce use of RTMPT transport when communicating with upstream hosts.

Reference: VPM Reference Information.

Flash streaming proxy supports policy enforcement based on RTMP traffic. The tables below lists the Flash-related VPM commands for generating policy. For more information about the VPM, refer to the *Visual Policy Manager Reference*.

The VPM objects related to Flash are listed below:

VPM Object	VPM Layer
Client Protocol (Column: service)	
User Agent (Column: source)	
Streaming Client (Column: service)	
Streaming Content Type (Column: service)	
Flash Streaming App Name (Column: destination)	Web Content
Protocol Methods (Column: service)	
Flash Streaming Stream Name (Column: destination)	Web Content
Request URL (Column: destination)	

Section H: Supported Streaming Media Clients and Protocols

This section describes the vendor-specific streaming protocols supported by the ProxySG.

Note: Blue Coat recommends upgrading to WMP version 9 or later. WMP versions 11 and higher do not support the Microsoft Media Services (MMS) protocol.

Supported Streaming Media Clients and Servers

The ProxySG supports Microsoft Windows Media, Flash Player, Apple QuickTime, and RealNetworks RealPlayer; however, the various players might experience unexpected behavior dependent upon certain SGOS configurations and features. Feature sections list such interactivities, as necessary. For a list of the most current versions of each supported client, refer to the Blue Coat SGOS *Release Notes* for this release.

Supported Flash Players and Servers

The Flash streaming proxy is compatible with current versions of Flash Media Server, client plug-ins, and browsers.

Supported Smooth Streaming Players and Servers

All servers and clients capable of Smooth Streaming are supported.

Supported Windows Media Players and Servers

The ProxySG supports the following versions and formats:

- ❑ Windows Media Player
- ❑ Windows Media Server
- ❑ Microsoft Silverlight

Note: Silverlight is supported when it streams Windows Media content from the WM server using WM-HTTP protocol. In this scenario, its interaction with the ProxySG appliance is similar to that of Windows Media Player, and, as such, is handled by the Windows Media proxy.

Supported Real Media Players and Servers

The ProxySG supports the following versions:

- ❑ RealOne Player
- ❑ RealPlayer
- ❑ RealServer
- ❑ Helix Universal Server

Note: Blue Coat recommends not deploying a Helix proxy between the ProxySG and a Helix server where the Helix proxy is the parent to the ProxySG. This causes errors with the Helix server. The reverse is acceptable (using a Helix proxy as a child to the ProxySG).

Supported QuickTime Players and Servers

The ProxySG supports the following versions, but in pass-through mode only:

- ❑ QuickTime Player
- ❑ Darwin Streaming Server
- ❑ Helix Universal Server

Supported Streaming Protocols

Each streaming media platform supports its own set of protocols. This section describes the protocols the ProxySG supports.

Flash Protocols

Flash streaming proxy supports the following RTMP-based protocols:

Supported Protocols	Supported Proxy Types	Features/Limitations
RTMP	Transparent	Full proxy feature support
RTMPT	Explicit, Transparent	Full proxy feature support
RTMPE	Transparent	Full proxy feature support of current RTMPE versions. Connections that use an unrecognized protocol version are passed through, without decryption or acceleration. For these connections, the Detail column in the Active Sessions report shows Encrypted, tunneled as unknown protocol version .
RTMPTE	Explicit, Transparent	Full proxy feature support of current RTMPTE versions. Connections that use an unrecognized protocol version are passed through, as described above.

Note: RTMP over SSL (RTMPS) is *not* currently supported.

Smooth Streaming Protocols

Smooth Streaming uses the HTTP protocol; the ProxySG supports Smooth Streaming over HTTP.

Windows Media Protocols

The ProxySG supports Windows Media content streamed over RTSP and HTTP. The following Windows Media transports are supported:

Client-side

- ❑ RTP over unicast UDP (RTSP over TCP, RTP over unicast UDP)
- ❑ Interleaved RTSP (RTSP over TCP, RTP over TCP on the same connection)
- ❑ RTP over multicast UDP (RTP over multicast UDP; for live content only)
- ❑ HTTP streaming
- ❑ MMS-UDP (Microsoft Media Streaming—User Data Protocol)
- ❑ MMS-TCP (Microsoft Media Streaming—Transmission Control Protocol)
- ❑ Multicast-UDP is the only delivery protocol supported for multicast. No TCP control connection exists for multicast delivery

Server-side

- ❑ Interleaved RTSP
- ❑ HTTP streaming
- ❑ MMS-TCP between the ProxySG and origin server for video-on-demand and live unicast content

Server-side RTP over UDP is not supported. If policy directs the RTSP proxy to use HTTP as server-side transport, the proxy denies the client request. The client then rolls over to MMS or HTTP.

Note: The MMS protocol is usually referred to as either MMS-TCP or MMS-UDP depending on whether TCP or UDP is used as the transport layer for sending streaming data packets. MMS-UDP uses a TCP connection for sending and receiving media control messages, and a UDP connection for streaming the actual media data. MMS-TCP uses TCP connections to send both control and data messages. The MMS protocol is not supported in WMP 11 and higher.

Real Media Protocols

The ProxySG supports the following Real Media protocols:

Client-Side

- ❑ HTTP streaming (RTSP and RDT over TCP tunneled through HTTP)—HTTP streaming is supported through a handoff process from HTTP to RTSP. HTTP accepts the connection and, based on the headers, hands off to RTSP. The headers identify an RTSP URL.

- ❑ RDT over unicast UDP (RTSP over TCP, RDT over unicast UDP)
- ❑ Interleaved RTSP (RTSP over TCP, RDT over TCP on the same connection)
- ❑ RDT over multicast UDP (RTSP over TCP, RDT over multicast UDP; for live content only)

Server-Side

- ❑ HTTP streaming
- ❑ Interleaved RTSP

Unsupported Protocols

The following Real Media protocols are not supported in this version of SGOS:

- ❑ PNA
- ❑ Server-side RDT/UDP (both unicast and multicast)

QuickTime Protocols

The ProxySG supports the following QuickTime protocols:

- ❑ HTTP streaming (RTSP and RDT over TCP tunneled through HTTP)—HTTP streaming is supported through a handoff process from HTTP to RTSP. HTTP accepts the connection and, based on the headers, hands off to RTSP. The headers identify an RTSP URL.
- ❑ RTP over unicast UDP (RTSP over TCP, RDT over unicast UDP)
- ❑ Interleaved RTSP (RTSP over TCP, RDT over TCP on the same connection)

Server-Side

- ❑ HTTP streaming
- ❑ Interleaved RTSP

Unsupported Protocols

The following QuickTime protocols are not supported in this version of SGOS:

- ❑ Server-side RTP/UDP, both unicast and multicast, is not supported.

Client-side multicast is not supported.

Chapter 25: Bandwidth Management

Bandwidth management (BWM) allows you to classify, control, and limit the amount of bandwidth used by different classes of network traffic flowing into or out of the ProxySG appliance. Network resource sharing (or link sharing) is accomplished by using a bandwidth-management hierarchy where multiple traffic classes share available bandwidth in a controlled manner.

Note: The ProxySG does not attempt to reserve any bandwidth on the network links that it is attached to or otherwise guarantee that the available bandwidth on the network can sustain any of the bandwidth limits which have been configured on it. The ProxySG can only shape the various traffic flows passing through it, and prioritize some flows over others according to its configuration.

By managing the bandwidth of specified classes of network traffic, you can accomplish the following:

- ❑ Guarantee that certain traffic classes receive a specified minimum amount of available bandwidth.
- ❑ Limit certain traffic classes to a specified maximum amount of bandwidth.
- ❑ Prioritize certain traffic classes to determine which classes have priority over available bandwidth.

Topics in this Section

This section includes information about the following topics:

- ❑ ["Bandwidth Management Overview"](#) on page 597
- ❑ ["Configuring Bandwidth Allocation"](#) on page 602
- ❑ ["Bandwidth Management Statistics"](#) on page 604
- ❑ ["Using Policy to Manage Bandwidth"](#) on page 606

Bandwidth Management Overview

To manage the bandwidth of different types of traffic that flow into, out of, or through the ProxySG, you must perform the following:

- ❑ Determine how many bandwidth classes you need and how to configure them to accomplish your bandwidth management goals. This includes determining the structure of one or more bandwidth hierarchies if you want to use priority levels to manage bandwidth.
- ❑ Create and configure bandwidth classes accordingly.

- ❑ Create policy rules using those bandwidth classes to identify and classify the traffic in the ProxySG.
- ❑ Enable bandwidth management.

Bandwidth management configuration consists of two areas:

- ❑ **Bandwidth allocation**—This is the process of creating and configuring bandwidth classes and placing them into a bandwidth class hierarchy. This process can be done using either the Management Console or the CLI. See ["Allocating Bandwidth"](#) on page 598.
- ❑ **Flow classification**—This is the process of classifying traffic flows into bandwidth management classes using policy rules. Policy rules can classify flows based on any criteria testable by policy. You can create policy rules using either the Visual Policy Manager (VPM), which is accessible through the Management Console, or by composing Content Policy Language (CPL). See ["Flow Classification"](#) on page 601.

Note: For more information about using VPM to create policy rules, refer to the *Visual Policy Manager Reference*. For information about composing CPL, refer to the *Content Policy Language Guide*.

Allocating Bandwidth

The process of defining bandwidth classes and grouping them into a bandwidth class hierarchy is called *bandwidth allocation*. Bandwidth allocation is based on:

- ❑ the placement of classes in a hierarchy (the parent/child relationships).
- ❑ the priority level of classes in the same hierarchy.
- ❑ the minimum and/or maximum bandwidth setting of each class.

For example deployment scenarios, see ["Bandwidth Allocation and VPM Examples"](#) on page 607.

Bandwidth Classes

To define a bandwidth class, you create the class, giving it a name meaningful to the purpose for which you are creating it. You can configure the class as you create it or edit it later. The available configuration settings are:

- ❑ **Parent:** Used to create a bandwidth-management hierarchy.
- ❑ **Minimum Bandwidth:** Minimum amount of bandwidth guaranteed for traffic in this class.
- ❑ **Maximum Bandwidth:** Maximum amount of bandwidth allowed for traffic in this class.
- ❑ **Priority:** Relative priority level among classes in the same hierarchy.

Parent Class

A parent class is a class that has children. When you create or configure a bandwidth class, you can specify another class to be its parent (the parent class must already exist). Both classes are now part of the same bandwidth-class hierarchy, and so are subject to the hierarchy rules (see "[Class Hierarchy Rules and Restrictions](#)" on page 600).

Minimum Bandwidth

Setting a minimum for a bandwidth class guarantees that class receives at least that amount of bandwidth, if the bandwidth is available. If multiple hierarchies are competing for the same available bandwidth, or if the available bandwidth is not enough to cover the minimum, bandwidth management is not be able to guarantee the minimums defined for each class.

Note: The ProxySG does not attempt to reserve any bandwidth on the network links that it is attached to or otherwise guarantee that the available bandwidth on the network can be used to satisfy bandwidth class minimums. The ProxySG can only shape the various traffic flows passing through it, and prioritize some flows over others according to its configuration.

Maximum Bandwidth

Setting a maximum for a bandwidth class puts a limit on how much bandwidth is available to that class. It does not matter how much bandwidth is available; a class can never receive more bandwidth than its maximum.

To prevent a bandwidth class from using more than its maximum, the ProxySG inserts delays before sending packets associated with that class until the bandwidth used is no more than the specified maximum. This results in queues of packets (one per class) waiting to be sent. These queues allow the ProxySG to use priority settings to determine which packet is sent next. If no maximum bandwidth is set, every packet is sent as soon as it arrives, so no queue is built and nothing can be prioritized.

Unlike minimums and priority levels, the maximum-bandwidth setting can purposely slow down traffic. Unused bandwidth can go to waste with the maximum-bandwidth setting, while the minimum-bandwidth settings and priority levels always distributes any unused bandwidth as long as classes request it. However, priority levels are not meaningful without a maximum somewhere in the hierarchy. If a hierarchy has no maximums, any class in the hierarchy can request and receive any amount of bandwidth regardless of its priority level.

Priority

When sharing excess bandwidth with classes in the same hierarchy, the class with the highest priority gets the first opportunity to use excess bandwidth. When the high-priority class uses all the bandwidth it needs or is allowed, the next class gets to use the bandwidth, if any remains. If two classes in the same hierarchy have the same priority, then excess bandwidth is shared in proportion to their maximum bandwidth setting.

Class Hierarchies

Bandwidth classes can be grouped together to form a class hierarchy. Creating a bandwidth *class* allows you to allocate a certain portion of the available bandwidth to a particular type of traffic. Putting that class into a bandwidth-class *hierarchy* with other bandwidth classes allows you to specify the relationship among various bandwidth classes for sharing available (unused) bandwidth.

The way bandwidth classes are grouped into the bandwidth hierarchy determines how they share available bandwidth among themselves. You create a hierarchy so that a set of traffic classes can share unused bandwidth. The hierarchy starts with a bandwidth class you create to be the top-level parent. Then you can create other bandwidth classes to be the children of the parent class, and those children can have children of their own.

To manage the bandwidth for any of these classes, some parent in the hierarchy must have a maximum bandwidth setting. The classes below that parent can then be configured with minimums and priority levels to determine how unused bandwidth is shared among them. If none of the higher level classes have a maximum bandwidth value set, then bandwidth flows from the parent to the child classes without limit. In that case, minimums and priority levels are meaningless, because all classes get all the bandwidth they need at all times. The bandwidth, in other words, is not being managed.

Class Hierarchy Rules and Restrictions

Certain rules and restrictions must be followed to create a valid BWM class hierarchy:

- ❑ Each traffic flow can only belong to one bandwidth management class.
You can classify multiple flows into the same bandwidth class, but any given flow is always counted as belonging to a single class. If multiple policy rules match a single flow and attempt to classify it into multiple bandwidth classes, the last classification done by policy applies.
- ❑ When a flow is classified as belonging to a bandwidth class, all packets belonging to that flow are counted against that bandwidth class.
- ❑ If a minimum bandwidth is configured for a parent class, it must be greater than or equal to the sum of the minimum bandwidths of its children.
- ❑ If a maximum bandwidth is configured for a parent class, it must be greater than or equal to the largest maximum bandwidth set on any of its children. It must also be greater than the sum of the minimum bandwidths of all of its children.
- ❑ The minimum bandwidth available to traffic directly classified to a parent class is equal to its assigned minimum bandwidth minus the minimum bandwidths of its children. For example, if a parent class has a minimum bandwidth of 600 kbps and each of its two children have minimums of 300 kbps, the minimum bandwidth available to traffic directly classified into the parent class is 0.

Relationship among Minimum, Maximum, and Priority Values

Maximum values can be used to manage bandwidth for classes whether or not they are placed into a hierarchy. This is not true for minimums and priorities, which can only manage bandwidth for classes that are placed into a hierarchy. Additionally, a hierarchy must have a maximum configured on a high-level parent class for the minimums and priorities to manage bandwidth.

This is because, without a maximum, bandwidth goes to classes without limit and there is no point to setting priorities or minimum guarantees. Bandwidth cannot be managed unless a maximum limit is set somewhere in the hierarchy.

When a hierarchy has a maximum on the top-level parent and minimums, maximums and priorities placed on the classes related to that parent, the following conditions apply:

- ❑ If classes in a hierarchy have minimums, the first thing that happens with available bandwidth is that all the minimum requests are satisfied. If the amount requested is less than the minimum for any class, it receives the entire amount, and its priority level does not matter.

Even though a minimum is considered to be a guaranteed amount of bandwidth, satisfying minimums is dependent on the parent being able to receive its own maximum, which is not guaranteed.

- ❑ When all of the classes in a hierarchy have had their minimums satisfied, any additional requests for bandwidth must be obtained. When a class requests more than its minimum, it must obtain bandwidth from its parent or one of its siblings. If, however, a class requests more than its maximum, that request is denied—no class with a specified maximum is ever allowed more than that amount.
- ❑ If a class does not have a minimum specified, it must obtain all of the bandwidth it requests from its parents or siblings, and it cannot receive any bandwidth unless all of the minimums specified in the other classes in its hierarchy are satisfied.
- ❑ Classes obtain bandwidth from their parents or siblings based on their priority levels—the highest priority class gets to obtain what it needs first, until either its entire requested bandwidth is satisfied or until it reaches its maximum. After that, the next highest priority class gets to obtain bandwidth, and this continues until either all the classes have obtained what they can or until the maximum bandwidth available to the parent has been reached. The amount available to the parent can sometimes be less than its maximum, because the parent must also participate in obtaining bandwidth in this way with its own siblings and/or parent if it is not a top-level class.


Flow Classification

You can classify flows to BWM classes by writing policy rules that specify the bandwidth class that a particular traffic flow belongs to. A typical transaction has four traffic flows:

1. Client inbound—Traffic flowing into the ProxySG from a client (the entity sending a request, such as a client at a remote office linked to the appliance).

2. Server outbound—Traffic flowing out of the ProxySG to a server.
3. Server inbound—Traffic flowing back into the appliance from a server (the entity responding to the request).
4. Client outbound—Traffic flowing back out of the appliance to a client.

The figure below shows the traffic flows between a client and server through the ProxySG.



Some types of traffic can flow in all four directions. The following example describes different scenarios that you might see with an HTTP request. A client sends a GET to the ProxySG (client inbound). The appliance then forwards this GET to a server (server outbound). The server responds to the ProxySG with the appropriate content (server inbound), and then the appliance delivers this content to the client (client outbound).

Policy allows you to configure different classes for each of the four traffic flows. See ["Using Policy to Manage Bandwidth"](#) on page 606 for information about classifying traffic flows with policy.

Configuring Bandwidth Allocation

You can use either the Management Console or the CLI to perform the following tasks:


- ❑ Enable or disable bandwidth management.
- ❑ Create and configure bandwidth classes.
- ❑ Delete bandwidth classes.
- ❑ View bandwidth management class configurations.

Note: If you plan to manage the bandwidth of streaming media protocols (Windows Media, Real Media, or QuickTime), Symantec suggests using the streaming features instead of the bandwidth management features described in this section. For information about the differences between these two methods, see ["Managing Streaming Media"](#) on page 529.


For conceptual information about bandwidth management, see ["Bandwidth Management Overview"](#) on page 597.

To create bandwidth classes and enable bandwidth management:

1. Select the **Configuration > Bandwidth Mgmt > BWM Classes > Bandwidth Classes** tab.



2. Click **New**. The Create Bandwidth Class dialog displays.
3. Create a new BWM class:
 - a. **Class name:** Assign a meaningful name for this class. The name can be up to 64 characters long; spaces are not allowed.
 - b. **Parent:** (Optional) To assign the class as a child of another parent class in the bandwidth class hierarchy, select an existing parent class from the drop-down list.
 - c. **Min. Bandwidth:** (Optional) Select **Min. Bandwidth** and enter a minimum bandwidth value in the field (kilobits per second (kbps)). The default minimum bandwidth setting is *unspecified*, meaning the class is not guaranteed a minimum amount of bandwidth.
 - d. **Max. Bandwidth:** (Optional) Select **Max. Bandwidth** and enter a maximum bandwidth value in the field. The default maximum bandwidth setting is *unlimited*, meaning the class is not limited to a maximum bandwidth value by this setting.
 - e. **Priority:** Select a priority level for this class from the **Priority** drop-down list—**0** is the lowest priority level and **7** is the highest. The default priority is **0**.
 - f. Click **OK** to close the dialog.



After you add a child class to a parent class, the parent class is denoted by a folder icon. Double-click the folder to view all of the child classes under that parent.

4. Select **Enable Bandwidth Management** (if not currently selected).
5. Click **Apply**.

To delete a BWM class:

Note: You cannot delete a class that is referenced by another class or by the currently installed policy. For instance, you cannot delete a class that is the parent of another class or one that is used in an installed policy rule. If you attempt to do so, a message displays explaining why this class cannot be deleted.

1. Select **Configuration > Bandwidth Management > BWM Classes > Bandwidth Classes**.
2. Highlight the class to delete and **Delete**.
3. Click **Yes** to delete the class.
4. Click **Apply**.

Bandwidth Management Statistics

The bandwidth management statistics tabs ("[Current Class Statistics](#)" and "[Total Class Statistics](#)") display the current packet rate and total number of packets served, the current bandwidth rate, and the total number of bytes served and packets dropped.

Current Class Statistics


The **Current Class Statistics** tab displays the following information for each bandwidth class:

- ❑ **Current Packet Rate:** current packets-per-second (pps) value.
- ❑ **Current Bandwidth:** current bandwidth in kilobits per second (Kbps).

To view current bandwidth management class statistics:

1. Select **Statistics > Bandwidth Mgmt. > Current Class Statistics**.

The high level bandwidth classes and their statistics are visible.



2. To view the statistics of child bandwidth classes, double-click the folder icon of the parent class.

The child classes become visible. A second double-click closes the folder.

See Also

- ❑ ["Using Policy to Manage Bandwidth"](#)

Total Class Statistics

The **Total Class Statistics** tab displays the following information for each bandwidth class:

- ❑ **Packets:** the total number of packets served.
- ❑ **Bytes:** the total number of bytes served.
- ❑ **Drops:** the total number of packets dropped.

To view total bandwidth management class statistics:

1. Select **Statistics > Bandwidth Management > Total Class Statistics**.

The high level bandwidth classes and their statistics are visible.

Current Class Statistics

Total Class Statistics

Bandwidth Classes

Bandwidth Class	Packets	Bytes	Drops
● bridging	0	0	0
● ftp	0	0	0
● http	0	0	0
● im	0	0	0
● streaming	0	0	0
● p2p	0	0	0
● service-info	0	0	0
● socks	0	0	0
● tcp-tunnel	0	0	0

Help

2. To view the statistics of child bandwidth classes, double-click the folder icon of the parent class. A second double-click closes the folder.

To clear bandwidth management statistics (CLI only):

1. To clear bandwidth management statistics for all bandwidth management classes, enter the following command in the CLI:

```
SGOS# clear-statistics bandwidth-management
```

2. To clear bandwidth management statistics for a particular class, enter the following command at the prompt:

```
SGOS# clear-statistics bandwidth-management class bandwidth_class_name
```

See Also

- ❑ ["Using Policy to Manage Bandwidth"](#)

Using Policy to Manage Bandwidth

After creating and configuring bandwidth management classes, create policy rules to classify traffic flows using those classes. Each policy rule can only apply to one of four traffic flow types:

- ❑ Client inbound
- ❑ Client outbound
- ❑ Server inbound
- ❑ Server outbound

You can use the same bandwidth management classes in different policy rules; one class can manage bandwidth for several types of flows based on different criteria. However, any given flow is always be counted as belonging to a single class. If multiple policy rules match a flow and try to classify it into multiple bandwidth classes, the last classification done by policy applies.

To manage the bandwidth classes you have created, you can either compose CPL or use the VPM. To see examples of policy using these methods, see "[Bandwidth Allocation and VPM Examples](#)" on page 607 or "[Policy Examples: CPL](#)" on page 614.

Bandwidth Allocation and VPM Examples


This section illustrates how to use the VPM to allocate bandwidth, arrange hierarchies, and create policy. It describes an example deployment scenario and the tasks an administrator must accomplish to manage the bandwidth for this deployment. For specific instructions about allocating bandwidth, see "[Configuring Bandwidth Allocation](#)" on page 602. For examples of CPL bandwidth management tasks, see "[Policy Examples: CPL](#)" on page 614.

Task One: Bandwidth Allocation

The administrator is responsible for managing the bandwidth of three branch offices. He was told to ensure that each office uses no more than half of its total link bandwidth for Web and FTP traffic. The total link bandwidth of each office is as follows:

- ❑ Office A: 1.5 Mb
- ❑ Office B: 1 Mb
- ❑ Office C: 2 Mb


He creates one bandwidth class for each of the three offices and configures the maximum bandwidth to an amount equal to half of the total link bandwidth of each, as shown below. He also creates policy rules for each class, as described in "Task One: VPM".




Each of the classes above has a maximum set at an amount equal to half of the total link bandwidth for each office. A hierarchy does not exist in this scenario.

Task One: VPM


The administrator has created one bandwidth class for each office, setting a maximum bandwidth on each one equal to the half of the total link bandwidth of each. Now he must create policy rules to classify the traffic flows.




The administrator launches the VPM and creates a new Web Access Layer, naming it **FTP/HTTP Limitations**. He selects the **Client IP Address/Subnet** object in the **Source** column, filling in the IP address and mask of the subnet used by **Office_A**.



He selects a **Combined Service Object** in the **Service** column, naming it **FTP/HTTP** and adding a **Client Protocol** for FTP and for HTTP.



He adds both protocols to the **At least one of these objects** field.




In the **Action** column, he selects **Manage Bandwidth**, naming it **Office_A** and setting it to manage the bandwidth of **Office_A** on the **Client side** in the **Outbound** direction.

He adds two more similar rules for the other two offices. He is able to reuse the same **Combined Service Object** in the **Service** column, but must add new objects specific to each office in the **Source** and **Action** columns. The order of the rules does not matter here, because each office, and thus each rule, is distinct because of its IP address/subnet mask configuration.

Task Two: Bandwidth Allocation

A few days later, the administrator gets a visit from the CEO of his company. She wants him to fix it so that she can visit any of the branch offices without having her own Web and FTP access slowed down unnecessarily.

The administrator creates two more classes for each office: one for the CEO and another for everyone else (employees). He sets the parent class of each new class to the appropriate class that he created in Task One. For example, he creates **Emp_A** and **CEO_A** and sets their parent class to **Office_A**. He also sets a priority level for each class: **0** (the lowest) for employees and **1** for the CEO. He then uses VPM to create additional policy rules for the new classes (see "[Task Two: VPM](#)" on page 610). This figure shows the hierarchical relationship among all of the classes.




The administrator now has three separate hierarchies. In each one, bandwidth is limited by the configuration of the parent class, and the two child classes are prioritized to determine how they share any unused bandwidth. Because no minimums have been set, the highest priority class has the first opportunity to use all of the available bandwidth; whatever is left then goes to the next priority class.


Priority levels are only effective among the classes in the same hierarchy. This means that the priority levels for the **Office_A** hierarchy do not affect the classes in the **Office_B** or **Office_C** hierarchies.

Task Two: VPM

Because the CEO wants to prioritize FTP and HTTP access among employees and herself, the administrator must create additional bandwidth classes (as described above in "Task Two: Bandwidth Allocation") and write policy rules to classify the traffic for the new classes.



He first edits each of the three VPM rules for the three offices. He edits each the Manage Bandwidth objects, changing the name of the objects to **Emp_A**, **Emp_B**, and **Emp_C** and changes the bandwidth class to the corresponding employee class.



Next, he creates three more rules for the CEO, moving them above the first three rules. For the CEO rules, he selects the same combined **FTP/HTTP** object in the **Service** column; in the **Action** column, he selects a **Manage Bandwidth** object configured for client side/outbound, as before, but this time, he names the objects **CEO_A**, **CEO_B**, and **CEO_C** and selects the corresponding CEO bandwidth class. In the **Source** column, he creates a **Combined Source Object**, naming it for the CEO. He combines the **Client IP/subnet** object already created for each office with a **User** object that he creates for the CEO.


The administrator places all three CEO rules above the employee rules, because the ProxySG looks for the first rule that matches a given situation and ignores the remaining rules. If he had placed the CEO rules below the employee rules, the appliance would never get to the CEO rules because the CEO's Web surfing client IP address matches both the CEO rules and the employee rules, and the ProxySG would stop looking after the first match. With the CEO rules placed first, the appliance applies the CEO rules to the CEO's Web surfing, and an employee's Web surfing does not trigger the CEO rules and instead skips ahead to the appropriate employee rule.

Task Three: Bandwidth Allocation

It soon becomes apparent that CEO visits are causing problems for the branch offices. At times, she uses all of the available bandwidth, resulting in decreased productivity throughout the office she visits. Also, management has complained that they have been given the same priority for FTP and HTTP traffic as regular employees, and they are requesting that they be given priority over employees for this type of traffic.

First, the administrator creates two new classes for each office. In this example, we look at the classes and configurations for the first office only. He creates a class called **Staff_A** and sets a minimum bandwidth of 500 kbps on it. He also creates a

class called **Mgmt_A**, setting the priority to 1 and the parent to **Staff_A**. He edits the class **Emp_A**, setting the parent to **Staff_A**. Finally, he edits the class **CEO_A**, changing the priority to 2. The resulting hierarchy is illustrated below. To see what the administrator did to the policy rules, see ["Task Three: VPM"](#) on page 612.



In the example illustrated above, employees and management combined are guaranteed a total of 500 kbps. The CEO's priority level has no effect until that minimum is satisfied. This means that the CEO can only use 250 kbps of bandwidth if the rest of the staff are using a total of 500 kbps. It also means that the CEO can use 750 kbps if no one else is using bandwidth at the time. In fact, any of the classes can use 750 kbps if the other classes use none.

Priority levels kick in after all of the minimums are satisfied. In this example, if the staff requests more than 500 kbps, they can only receive it if the CEO is using less than 250 kbps. Now notice that the minimum setting for the staff is set on the parent class, **Staff_A**, and not on the child classes, **Emp_A** or **Mgmt_A**. This means that the two child classes, representing employees and management, share a minimum of 500 kbps. But they share it based on their priority levels. This means that management has priority over employees. The employees are only guaranteed a minimum if management is using less than 500 kbps.


Task Three: VPM

The administrator has added additional classes for each office and edited the existing employee classes, as described above in ["Task Three: Bandwidth Allocation"](#) on page 611. One of the new classes he added for each office is a parent class that does not have traffic classified to it; it was created to provide a minimum amount of bandwidth to its child classes. Not every class in the hierarchy has to have a traffic flow. This means that he needs to add just three more rules for the three new management classes. For the management rules, he selects the same combined **FTP/HTTP** object in the **Service** column; in the **Action** column, he selects a **Manage Bandwidth** object configured for client side/outbound with the bandwidth class one of the management classes (**Mgmt_A**, **Mgmt_B**, or **Mgmt_C**). In the **Source** column, he creates a **Combined Source Object** containing the subnet object for the office and the **Group** object for management.

The management rules must go above the employee rules, although it does not matter where they are placed in relation to the CEO rules. This would not be true if the CEO was part of the same group as management, however. If that were true, the CEO rules would still need to go on top.

Task Four: Bandwidth Allocation

The administrator decided later that he needed to guarantee employees some bandwidth. He configures a minimum for the class **Emp_A**, as illustrated below.



He decides to leave the minimum on the parent class **Staff_A** and not to set a minimum for the class **Mgmt_A**. This is okay, because the minimum of the parent class is available to its children if the parent class does not use all of it, and the only way that the CEO can get more than 250 kbps is if the employees and management combined use less than 500.

This last change does not require additional changes to policy; the administrator has added a minimum to a class that he has already classified for traffic using policy.

In the above scenario, the class called **Staff_A** does not have traffic configured for it—it was created to guarantee bandwidth minimums for its child classes. However, if it were configured for traffic, it would have a practical minimum of 300 kbps. The practical minimum of a parent class is equal to its assigned minimum bandwidth minus the minimums of its children. In that case, if the parent class **Staff_A** used 300 kbps and the child class **Emp_A** used 200 kbps, the child class **Mgmt_A** would not receive any bandwidth unless the class **CEO_A** was using less than 250 kbps. Under those circumstances, the administrator probably also needs to create a minimum for management.

Task Five: Bandwidth Allocation

The CEO makes another request, this time for the main office, the one the administrator himself works from. This office uses the content filtering feature of the ProxySG to control the types of Web sites that employees are allowed to view. Although the office uses content filtering, access to sports sites is not restricted because the CEO is a big fan.

The administrator creates a bandwidth management class called **Sports** with a maximum bandwidth of 500 kbps and launches VPM to create policy for this class as described below.

Task Five: VPM

To classify traffic for the **Sports** class, the administrator opens VPM, creates a Web Access Layer, and sets the **Destination** column to the **Category** object that includes sports viewing (content filtering is already set up in VPM). He sets the **Action**

column to the **Manage Bandwidth** object, selecting **Server side/Inbound** and the **Sports** bandwidth class he created. After installing the policy and verifying that bandwidth management is enabled, he is finished.

Policy Examples: CPL

The examples below are complete in themselves. The administrator uses CLI to create and configure bandwidth management classes and writes CPL to classify traffic flow for these classes. These examples do not make use of a bandwidth class hierarchy. For examples of hierarchies, see "[Bandwidth Allocation and VPM Examples](#)" on page 607.

Example One: CPL

In this example, the administrator of a college is asked to prevent college students from downloading MP3 files during peak hours, while still allowing the music department to download MP3 files at any time. The CPL triggers used are authentication and/or source subnet and MIME type. The action taken is to limit the total amount of bandwidth consumed by students to 40 kbps.

CLI commands:

```
SGOS#(config) bandwidth-management
SGOS#(config bandwidth-management) create mp3
SGOS#(config bandwidth-management) edit mp3
SGOS#(config bw-class mp3) max-bandwidth 40
```

CPL:

```
define condition student_mp3_weekday
    client_address=student_subnet response_header.Content-Type="audio/
mpeg" \
    weekday=1..5 hour=9..16
end condition

<proxy>
    condition=student_mp3_weekday limit_bandwidth.server.inbound(mp3)
```

Example Two: CPL

In this example, an administrator must restrict the amount of bandwidth used by HTTP POST requests for file uploads from clients to 2 Mbps. The CPL trigger used is request method, and the action taken is to throttle (limit) the amount of bandwidth used by client side posts by limiting inbound client side flows.

CLI:

```
SGOS#(config) bandwidth-management
bandwidth-management) create http_post
SGOS#(config bandwidth-management) edit http_post
SGOS#(config bw-class http_post) max-bandwidth 2000
```

CPL:

```
define condition http_posts
    http.method=POST
end condition

<proxy>
    condition=http_posts limit_bandwidth.client.inbound(http_post)
```

Example Three: CPL

In this example, the administrator of a remote site wants to limit the amount of bandwidth used to pre-populate the content from headquarters to 50 kbps during work hours. The CPL triggers used are current-time and pre-population transactions. The action taken is to limit the total amount of bandwidth consumed by pre-pop flows.

CLI:

```
SGOS#(config) bandwidth-management
SGOS#(config bandwidth-management) create pre-pop
SGOS#(config bandwidth-management) edit pre-pop
SGOS#(config bw-class pre-pop) max-bandwidth 50
```

CPL:

```
define condition prepop_weekday
    content_management=yes weekday=1..5 hour=9..16
end condition

<proxy>
    condition=prepop_weekday limit_bandwidth.server.inbound(pre-pop)
```


Chapter 26: Configuring Access Logging

Access logging allows you to track Web usage for the entire network or specific information on user or department usage patterns. These logs and reports can be made available in real-time or on a scheduled basis. This chapter describes access logging and provides procedures for enabling access logging and configuring upload schedules.

Note: *Event logging* is not the same as access logging. Event logging allows you to specify the types of system events logged, the size of the event log, and to configure Syslog monitoring.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ ["About Access Logging"](#) on page 617
- ❑ ["Enabling or Disabling Access Logging"](#) on page 619
- ❑ ["Configuring a Log for Uploading"](#) on page 620
- ❑ ["Testing Access Log Uploading"](#) on page 622
- ❑ ["Viewing Access-Log Statistics"](#) on page 623
- ❑ ["Example: Using VPM to Prevent Logging of Entries Matching a Source IP"](#) on page 626

About Access Logging

SGOS can create access logs for the traffic flowing through the system; in fact, each protocol can create an access log record at the end of each transaction for that protocol (such as for each HTTP request).

Note: The only data that can be logged in an access log on the ProxySG are the access-log fields and the CPL fields (found in [Chapter 30: "Access Log Formats"](#) on page 657).


These log records can be directed to one or more log *facilities*, which associates the logs with their configured log formats, upload schedules, and other customizable components. In addition, access logs can be encrypted and digitally signed before uploading.

Data stored in log facilities can be automatically uploaded to a remote location for analysis and archive purposes. The uploads can take place using HTTP, FTP, or one of several proprietary protocols. After they are uploaded, reporting tools such as Blue Coat Reporter can be used to analyze the log files. For information on using Blue Coat Reporter, refer to the *Blue Coat Reporter Initial Configuration Guide*.

About Facilities

A log facility is a separate log that contains a single logical file and supports a single log format. The facility contains the file's configuration and upload schedule information as well as other configurable information such as how often to rotate (switch to a new log) the logs at the destination, any passwords needed, and the point at which the facility can be uploaded.

Multiple access log facilities are supported, although each access log supports a single log format. You can log a single transaction to multiple log facilities through a global configuration setting for the protocol that can be modified on a per-transaction basis through policy.



Access Logging Protocols and Formats

The following protocols support configurable access logging:

- ❑ CIFS
- ❑ Endpoint Mapper
- ❑ FTP
- ❑ HTTP
- ❑ HTTPS Forward Proxy
- ❑ HTTPS Reverse Proxy

- ❑ Instant Messaging
- ❑ Peer-to-peer (P2P)
- ❑ RealMedia/QuickTime
- ❑ SOCKS
- ❑ SSL
- ❑ TCP Tunnel
- ❑ Telnet
- ❑ Windows Media

SGOS can create access logs with any one of a number of log formats, and you can create additional types using custom or ELFF format strings. The log types supported are:

- ❑ NCSA common log format
- ❑ SQUID-compatible format
- ❑ ELFF (W3C Extended Log File Format)
- ❑ Custom, using the strings you enter
- ❑ SurfControl, a log format compatible with the SurfControl Reporter tool

The log facilities, each containing a single logical file and supporting a single log format, are managed by policy (created through the Visual Policy Manager (VPM) or Content Policy Language (CPL)), which specifies the destination log format and log file.

Enabling or Disabling Access Logging


You can globally enable or disable access logging. If access logging is disabled, logging is turned off for all log objects, even if logging policy exists or logging configurations are set.

After globally enabled, connection information is sent to the default log facility for the service. For example, HTTP traffic is logged to the main file.

By default, access logging is disabled on all new systems, but certain protocols are configured to use specific logs by default. When access logging is enabled, logging begins immediately for all configured protocols.

To enable or disable access logging:

1. Select **Configuration > Access Logging > General > Default Logging**.



2. Select **Enable** to enable access logging or deselect it to disable access logging.
3. Click **Apply**.

Configuring a Log for Uploading

The *upload schedule* defines the frequency of the access logging upload to a remote server, the time between connection attempts, the time between keep-alive packets, the time at which the access log is uploaded, and the protocol that is used. When configuring an upload schedule, you can specify either *periodic uploading* or *continuous uploading*. Both periodic and continuous uploading can send log information from an ProxySG appliance farm to a single log analysis tool. This allows you to treat multiple appliances as a single entity and to review combined information from a single log file or series of related log files.

With periodic uploading, the SGOS software transmits log entries on a scheduled basis (for example, once daily or at specified intervals) as entries are batched, saved to disk, and uploaded to a remote server.

Note: When you configure a log for continuous uploading, it continues to upload until you stop it. To stop continuous uploading, switch to periodic uploading temporarily. This is sometimes required for gzip or encrypted files, which must stop uploading before you can view them.

With continuous uploading, the ProxySG continuously *streams* new access log entries from the device memory to a remote server. Here, *streaming* refers to the real-time transmission of access log information. The SGOS software transmits access log entries using the specified client, such as FTP client. A keep-alive is sent to keep the data connection open.

Continuous uploading allows you to view the latest logging information almost immediately, send log information to a log analysis tool for real-time processing and reporting, maintain the ProxySG performance by sending log information to a remote server (avoiding disk writes), and save device disk space by saving log information on the remote server.

If the remote server is unavailable to receive continuous upload log entries, the SGOS software saves the log information on the device disk. When the remote server is available again, the appliance resumes continuous uploading.

Note: If you do not need to analyze the upload entries in real time, use periodic uploading because it is more reliable than continuous uploading.

If there is a problem configuring continuous uploading to Microsoft Internet Information Server (IIS), use periodic uploading instead.

To configure the upload schedule:

1. Select **Configuration > Access Logging > Logs > Upload Schedule**.

The screenshot shows the 'Upload Schedule' configuration page. At the top, there are tabs: 'Logs', 'General Settings', 'Upload Client', and 'Upload Schedule'. The 'Upload Schedule' tab is active. Below the tabs, there is a 'Log' dropdown menu set to 'main'. Below that, there is a section for 'Upload type:' with two radio buttons: 'continuously' (selected) and 'periodically'. Below the radio buttons, there are two input fields: 'Wait between connect attempts:' set to '60' seconds, and 'Time between keep-alive log packets:' set to '300' seconds. Below these, there is a section for 'Rotate the log file:' with two radio buttons: 'Daily at' (selected) and 'Every'. The 'Daily at' option has a dropdown menu set to '2:00 a.m.'. The 'Every' option has input fields for '8' hours and '0' minutes. To the right of the 'Rotate the log file:' section, there is a 'Rotate Now' button. Numbered callouts point to various elements: 2 points to the 'Log' dropdown, 3a points to the 'continuously' radio button, 3b points to the 'Wait between connect attempts' input field, 3c points to the 'Time between keep-alive log packets' input field, 4a points to the 'Daily at' radio button, 4b points to the 'Every' radio button, and 5 points to the 'Rotate Now' button.

2. From the **Log** drop-down list, select the log type.
3. Select the **Upload Type**:
 - a. Select **continuously** (stream access log entries to a remote server) or **periodically** (transmit on a scheduled basis).
 - b. To change the time between connection attempts, enter the new time (in seconds) in the **Wait between connect attempts** field.

- c. (Only accessible if you are updating continuously) To change the time between keep-alive packets, enter the new time (in seconds) in the **Time between keep-alive log packets** field.

Keepalives maintain the connection during low periods of system usage. When no logging information is being uploaded, the SGOS software sends a keep-alive packet to the remote server at the interval you specify, from 1 to 65535 seconds. If you set this to 0 (zero), you effectively disable the connection during low usage periods. The next time that access log information needs to be uploaded, the ProxySG automatically reestablishes the connection.

4. Determine when logs are uploaded or rotated:
 - a. (Optional) From the **Daily at** drop-down list, specify the time of day to log update (for periodic uploads) or rotate (for continuous uploads).
 - b. (Optional) To have the log uploaded or rotated on a daily basis, select **Every** and enter the time between uploads.
5. **Rotate or Upload Now:**
 - Continuous Upload: *Log rotation* helps prevent logs from growing excessively large. Especially with a busy site, logs can grow quickly and become too big for easy analysis. With log rotation, the SGOS software periodically creates a new log file, and archives the older one without disturbing the current log file.
 - Periodic Upload: You can upload the access logs now or you can cancel any access-log upload currently in progress (if you are doing periodic uploads). You can rotate the access logs now (if you are doing continuous uploads). These actions do not affect the next scheduled upload time.
 - **Cancel upload** (for periodic uploads) allows you to stop repeated upload attempts if the Web server becomes unreachable while an upload is in progress. Clicking this sets log uploading back to idle if the log is waiting to retry the upload. If the log file is in the process of uploading, it takes time for it to take effect.
6. Click **Apply**.

Testing Access Log Uploading

For the duration of the test, configure the event log to use the verbose event level (see "[Selecting Which Events to View](#)" on page 1309). This logs more complete log information. After you test uploading, you can check the event log for the test upload event and determine whether any errors occurred (go to **Statistics > Event Logging**). You cannot check the event log.

To test access log uploading:

You can do a test access log upload. Before you begin, make sure you have configured the upload client completely.

1. Select **Configuration > Access Logging > Logs > Upload Client**.
2. Click **Test Upload**.

3. Click **OK** in the Test upload dialog.
4. Check the event log for upload results: go to **Statistics > Event Logging**.

Viewing Access-Log Statistics

You can view some access log statistics by navigating to **Statistics > Advanced** and clicking **Access Log**. Statistics you can view from **Statistics > Advanced** include:

- ❑ **Show list of all logs:** The access log manages multiple log objects internally. These are put together as one logical access log file when the file is uploaded.

The show list shows the available internal log objects for easy access. To download part of the access log instead of the whole log file, click on the individual log object shown in the list. The latest log object can be identified by its timestamp.


Note: If you have multiple access logs, each access log has its own list of objects.

- ❑ **Show access log statistics:** The statistics of an individual access log is shown.
- ❑ **Show statistics of all logs:** The statistics of all the access logs on the system are displayed in a single list.
- ❑ **Show last N bytes in the log:** The last *N* bytes in the log are shown.
- ❑ **Show last part of log every time it changes:** A stream of the latest log entries is shown on the page as they are written in the system.
- ❑ **Show access log tail with optional refresh time:** A refresh from the browser displays the latest log entries.
- ❑ **Show access log objects:** The statistics of individual access log objects are displayed.
- ❑ **Show all access log objects:** The statistics of all access log object are displayed in a single list.

Viewing the Access Log Tail

To display the access log tail:

1. Select **Statistics > Access Logging > Log Tail**.



2. From the **Log** drop-down list, select the log to view.
3. Click **Start Tail** to display the access log tail.
The ProxySG displays a maximum of 500 lines. Entries that pre-date these 500 lines are not displayed.
4. Click **Stop Tail** to stop the display or **Clear Tail** to clear the display.

Viewing the Log File Size

The **Log Size** tab displays current log statistics:

- ❑ Whether the log is being uploaded (Table 26–1, "Log Writing Status Description" describes upload statuses)
- ❑ The current size of all access log objects
- ❑ Disk space usage
- ❑ Last modified time
- ❑ Estimated size of the access log file, once uploaded

Table 26–1 Log Writing Status Description

Status	Description
active	Log writing is active.
active - early upload	The early upload threshold has been reached.
disabled	An administrator has disabled logging.
idle	Log writing is idle.
initializing	The system is initializing.
shutdown	The system is shutting down.

Table 26–1 Log Writing Status Description (Continued)

stopped	The access log is full. The maximum log size has been reached.
unknown	A system error has occurred.

Estimated compressed size of the uploaded access log and ProxySG access log size might differ during uploading. This occurs because new entries are created during the log upload.

To view the access log size statistic:

1. Select **Statistics > Access Logging > Log Size**.

2. From the **Log** drop-down list, select a log to view.

Viewing Access Logging Status

The SGOS software displays the current access logging status on the Management Console. This includes separate status information about:

- ❑ The writing of access log information to disk
- ❑ The client the ProxySG uses to upload access log information to the remote server

To view access logging upload status:

1. Select **Statistics > Access Logging > Upload Status**.

2. Under **Status of Last Upload**, check the appropriate status information displayed in the **Upload client** field.

3. Check the other status information. For information about the status, see the table below.

Table 26–2 Upload Status Information


Status	Description
Connect time	The last time a client connection was made or attempted.
Remote filename	The most recent upload filename. If an access log was encrypted, only the encrypted access log file (the ENC file) displays.
Remote size	The current size of the upload file. If an access log was encrypted, only the encrypted access log file size (the ENC file) displays. The private key file (the DER file) varies, but is usually about 1 Kb.
Maximum bandwidth	The maximum bandwidth used in the current or last connection.
Current bandwidth	The bandwidth used in the last second (available only if currently connected).
Final result	The result of the last upload attempt (success or failure). This is available only if not connected.

Example: Using VPM to Prevent Logging of Entries Matching a Source IP


Complete the following steps to prevent a source IP address from being logged.

To prevent a source IP address from being logged:


1. Create a Web Access Layer:
 - a. Select **Configuration > Policy > Visual Policy Manager**; click **Launch**.



- b. In the VPM, select **Policy > Add Web Access Layer**.
- c. Enter a layer name into the dialog that appears and click **OK**.



2. Add a **Source** object:
 - a. Right click on the item in the **Source** column; select **Set**.
 - b. Click **New**; select **Client IP Address/Subnet**.
3. Enter an IP address or Subnet Mask in the dialog that appears and click **Add**; click **Close** (or add additional addresses and then click **Close**); click **OK**.
4. Add an Action object to this rule:
 - a. Right-click on the item in the **Action** column; select **Set**.
 - b. Click **New** in the Set Action Object dialog that appears; select **Modify Access Logging**.



- c. To disable a particular log, click **Disable logging to** and select that log from the drop-down list; to disable all access logging, click **Disable all access logging**.
5. Click **OK**; click **OK** again; close the VPM window and click **Yes** in the dialog to save your changes.

Chapter 27: Configuring the Upload Client

The ProxySG supports three types of upload client:

- ❑ FTP client, the default
- ❑ HTTP client
- ❑ Custom client
- ❑ Symantec Reporter client

Symantec also supports secure FTP (FTPS), secure HTTP (HTTPS), secure and Custom client.

The Custom client can be used for special circumstances, such as supporting Syslog TCP. Custom client is based on plain sockets.

Note: You must have a socket server to use the Custom client.

Topics in this Chapter:

This chapter includes information about the following topics:

- ❑ ["Encrypting the Access Log" on page 630](#)
- ❑ ["Importing an External Certificate" on page 630](#)
- ❑ ["Digitally Signing Access Logs" on page 631](#)
- ❑ ["Disabling Log Uploads" on page 634](#)
- ❑ ["Decrypting an Encrypted Access Log" on page 635](#)
- ❑ ["Verifying a Digital Signature" on page 635](#)
- ❑ ["Editing Upload Clients" on page 635](#)

The general options you enter in the **Upload Client** tab affect all clients. Specific options that affect individual clients are discussed in the FTP client, HTTP client, or Custom client, or the `access-log ftp-client, https-client, or custom-client` CLI commands.

Only one client can be used at any one time. All four can be configured, but only the selected client is used.

The SGOS software provides access logging with two types of uploads to a remote server:

- ❑ Continuous uploading, where the device continuously streams new access log entries from the device memory to a remote server.
- ❑ Scheduled (periodic) uploading, where the device transmits log entries on a scheduled basis. See ["Configuring Access Logging"](#) for more information.

The SGOS software allows you to upload either compressed access logs or plain-text access logs. The device uses the gzip format to compress access logs. Gzip-compressed files allow more log entries to be stored in the device. Advantages of using file compression include:

- ❑ Reduces the time and resources used to produce a log file because fewer disk writes are required for each megabyte of log-entry text.
- ❑ Uses less bandwidth when the device sends access logs to an upload server.
- ❑ Requires less disk space.

Compressed log files have the extension `.log.gz`. Text log files have the extension `.log`.

For greater security, you can configure the SGOS software to:

- ❑ Encrypt the access log
- ❑ Sign the access log

Encrypting the Access Log

To encrypt access log files, you must first place an external certificate on the ProxySG (see ["Importing an External Certificate"](#) on page 630). The device derives a session key from the public key in the external certificate and uses it to encrypt the log. When an access log is encrypted, two access log files are produced: an ENC file (extension `.enc`), which is the encrypted access log file, and a DER file (extension `.der`), which contains the ProxySG session key and other information. You need four things to decrypt an encrypted access log:

- ❑ The ENC file
- ❑ The DER file
- ❑ The external (public key) certificate
- ❑ The corresponding private key

For information about decrypting a log, see ["Decrypting an Encrypted Access Log"](#) on page 635.


Note: The encryption feature is not available for custom clients.

Importing an External Certificate

You can import an X.509 certificate into the ProxySG to use for encrypting data.

To Import an external certificate:

1. Copy the certificate onto the clipboard.
2. Select **Configuration > SSL > External Certificates**.
3. Click **Import**.



4. Enter the name of the external certificate into the **External Cert Name** field and paste the certificate into the **External Certificate** field. Be sure to include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` statements.
5. Click **OK**.
6. Click **Apply** to commit the changes to the ProxySG.

Deleting an External Certificate

To delete an external certificate:

1. Select **Configuration > SSL > External Certificates**.
2. Highlight the name of the external certificate to be deleted.
3. Click **Delete**.
4. Click **OK** in the Confirm Delete dialog that displays.
5. Click **Apply**.

Digitally Signing Access Logs

You can digitally sign access logs to certify that a particular ProxySG wrote and uploaded this log file. Signing is supported for both content types—text and gzip—and for both upload types—continuous and periodic. Each log file has a signature file associated with it that contains the certificate and the digital

signature for verifying the log file. The signature file has the same name as the access log file but with a `.sig` extension; that is, `filename.log.sig`, if the access log is a text file, or `filename.log.gzip.sig`, if the access log is a gzip file.

Note: Signing is disabled by default.

See one of the following topics for more information:

- ❑ ["Introduction to Digitally Signing Access Logs"](#)
- ❑ ["Configuring the Upload Client to Digitally Sign Access Logs"](#) on page 632

Introduction to Digitally Signing Access Logs

You can digitally sign your access log files with or without encryption. If the log is both signed and encrypted, the signing operation is done first, meaning that the signature is calculated on the unencrypted version of the file. You must decrypt the log file before verifying the file. Attempting to verify an encrypted file fails.

When you create a signing keyring (which must be done before you enable digital signing), keep in mind the following:

- ❑ The keyring must include an external certificate. (An external certificate is one for which the ProxySG does not have the private key.)
- ❑ The certificate purpose must be set for **smime** signing. If the certificate purpose is set to anything else, you cannot use the certificate for signing.
- ❑ Add the `%c` parameter in the filenames format string to identify the keyring used for signing. If encryption is enabled along with signing, the `%c` parameter expands to `keyringName_Certname`.

Note: The signing feature is not available for custom clients.

For information about verifying a log, see ["Verifying a Digital Signature"](#) on page 635.

Continue with ["Configuring the Upload Client to Digitally Sign Access Logs"](#).

Configuring the Upload Client to Digitally Sign Access Logs

This section discusses how to configure the upload client to digitally sign access logs. For more information, see ["Introduction to Digitally Signing Access Logs"](#) on page 632.

To configure the upload client:

1. Select **Configuration > Access Logging > Logs > Upload Client**.

2. From the **Log** drop-down list, select the log facility to configure. The facility must exist before it displays in this list.

3. Select and configure the client type:

- a. From the **Client type** drop-down list, select the upload client to use. Only one client can be configured for each log facility.
- b. Click **Settings** to customize the upload client.

For information on customizing the clients, skip to ["Editing the FTP Client"](#) on page 636, ["Editing the HTTP Client"](#) on page 637, ["Editing the Custom client"](#) on page 639, or ["Editing the Custom SurfControl Client"](#) on page 640.

For information about testing the upload client, see ["Testing Access Log Uploading"](#) on page 622.

4. Configure **Transmission Parameters**, if applicable:

- a. (Optional) To use an external certificate to encrypt the uploaded log facility, select an external certificate from the **Encryption Certificate** drop-down list. You must first import the external certificate to the ProxySG appliance (see ["Importing an External Certificate"](#) on page 630).

The encryption option is not available for Custom clients.

- b. (Optional) To enable the digital signature of the uploaded access log, select a keyring from the **Keyring Signing** drop-down list. The signing keyring, with a certificate set to **smime**, must already exist. A certificate set to any other purpose cannot be used for digital signatures.

The digital signing option is not available for Custom clients.

- c. Select one of the **Save the log file as** radio buttons to determine whether the access log that is uploaded is compressed (**gzip file**, the default) or not (**text file**).

Note: If you are configuring a SurfControl Custom clientCustom client, select the **text file** radio button.

If you select **text file**, you can change the **Send partial buffer after *n* seconds** field to the time you need (30 seconds is the default).

This field configures the maximum time between text log packets, meaning that it forces a text upload after the specified length of time even if the internal log buffer is not full. If the buffer fills up before the time specified in this setting, the text uploads right away, and is not affected by this maximum setting.

Note: If you selected **gzip file**, the **Send partial buffer after *n* seconds** field is not configurable. Also, this setting is only valid for continuous uploading (see ["Configuring Access Logging"](#) on page 617 for information about continuous uploading).

- d. (Optional) To manage the bandwidth for this log facility, select a bandwidth class from the **Bandwidth Class** drop-down list.

The default setting is **none**, which means that bandwidth management is disabled for this log facility by default.

Note: Before you can manage the bandwidth for this log facility, you must first create a bandwidth-management class. It is the log facility that is bandwidth-managed—the upload client type does not affect this setting. See ["Bandwidth Management"](#) on page 597 for information about enabling bandwidth management and creating and configuring the bandwidth class. Less bandwidth slows down the upload, while more could flood the network.

5. Click **Apply**.

See Also

["Verifying a Digital Signature"](#) on page 635

["Digitally Signing Access Logs"](#) on page 631

Disabling Log Uploads

To disable log uploads, set the upload client-type to none.

To disable an upload:

1. Select **Configuration > Access Logging > Logs > Upload Client**.
2. Select the log facility for which you want to disable an upload from the **Log** drop-down menu.

3. Select **NONE** from the **Client type** drop-down menu.
4. Click **Apply**.

Decrypting an Encrypted Access Log

To decrypt an encrypted access log, you must concatenate the DER and ENC files (with the DER file in front of the ENC file) and use a program such as OpenSSL for decryption. For example, use the following UNIX command and a tool such as OpenSSL to concatenate the DER and ENC files and decrypt the resulting file:

```
cat path/filename_of_DER_file path/filename_of_ENC_file | openssl
smime -decrypt -inform DER -binary -inkey path/filename_of_private_key
-recipient path/filename_of_external_certificate -out path/
filename_for_decrypted_log_file
```

You can also download a script based on the OpenSSL tool for decryption. Go to https://download.bluecoat.com/release/SG4/files/accesslog_decrypt.zip.

Verifying a Digital Signature

If the file whose digital signature you want to verify is also encrypted, you must decrypt the file prior to verifying the signature. (See "Decrypting an Encrypted Access Log" on page 635 above for more information.)

You can use a program such as OpenSSL to verify the signature. For example, use the following command in OpenSSL:

```
openssl smime -CAfile cacrt -verify -in filename.sig -content
filename.log -inform DER -out logFile
```

where

<i>cacrt</i>	The CA certificate used to issue the certificate in the signature file.
<i>filename.sig</i>	The file containing the digital signature of the log file.
<i>filename.log</i>	The log file generated after decryption. If the access log is a gzip file, it contains a .gz extension.
<i>logFile</i>	The filename that is generated after signature verification.

Editing Upload Clients

Four upload clients are supported by Blue Coat: FTP, HTTP, and Custom. Each of these clients are described below. You can also create a SurfControl upload client.

Multiple upload clients can be configured per log facility, but only one can be enabled and used per upload.

Editing the FTP Client

To edit the FTP client:

1. Select **Configuration > Access Logging > Logs > Upload Client**.
2. Select **FTP Client** from the **Client type** drop-down list. Click the **Settings** button.

3. Select the primary or alternate FTP server to configure from the **Settings for** drop-down list.
4. Fill in the server fields, as appropriate:
 - a. **Host:** The name of the upload client host. If the **Use secure connections (SSL)** check box is selected, the host name must match the host name in the certificate presented by the server. The host can be defined as an IPv4 or IPv6 address, or a domain name that resolves to an IPv4 or IPv6 address.
 - b. **Port:** If an IP address is entered for the host, specify a port number; the default is 21 for FTP clients.
 - c. **Path:** The directory path where the access log is uploaded on the server.
 - d. **Username:** This is the username that is known on the host you are configuring.
 - e. **Change Password:** Change the password on the FTP; the Change Password dialog displays; enter and confirm the new password; click **OK**.
5. **Filename:** The **Filename** field is comprised of text and/or specifiers. The default filename includes specifiers and text that indicate the log name (%f), name of the external certificate used for encryption, if any (%c), the fourth parameter of the ProxySG IP address (%l), the date and time (Month: %m, Day: %d, Hour: %H, Minute: %M, Second: %S), and the .log or .gzip.log file extension.

Note: Be cautious if you change the **Filename** field. If an ongoing series of access logs files are produced and you do not have time-specifiers in this field, each access log file produced overwrites the old file. Also, if you use more than one external certificate to encrypt logs, include the %c specifier in the **Filename** field to keep track of which external certificate was used to encrypt the uploaded log file.

6. **Secure Connections:** If you use FTPS, select the **Use secure connections (SSL)** check box. The remote FTP server must support FTPS.
7. **Local Time:** If you want the upload to reflect the local time it was uploaded instead of Universal Time Coordinates (UTC), select **Local Time**.
8. **Use PASV:** With **Use PASV** selected (the default), the ProxySG connects to the FTP server. With **Use PASV** de-selected, the FTP server uses the PORT command to connect to the ProxySG.
9. Click **OK**.
10. Click **Apply**.

Editing the HTTP Client

Access log uploads done through an HTTP/HTTPS client use the HTTP PUT method. The destination HTTP server (where the access logs are being uploaded) must support this method. Microsoft's IIS allows the server to be directly configured for write (PUT/DELETE) access. Other servers, such as Apache, require installing a new module for the PUT method for access log client uploads. You can create either an HTTP or an HTTPS upload client through the HTTP Client dialog. (Create an HTTPS client by selecting **Use secure connections (SSL)**.)

Note: To create an HTTPS client, you must also import the appropriate CA Certificate. For more information, see "[Importing CA Certificates](#)" on page 1141.

To edit the HTTP client:

1. Select **Configuration > Access Logging > Logs > Upload Client**.
See [Chapter 27: "Configuring the Upload Client"](#) on page 629 for configuration information.
2. Select **HTTP Client** from the **Client type** drop-down list. Click **Settings**.

3. From the **Settings for** drop-down list, select the primary or alternate HTTP server to configure.
4. Fill in the server fields, as appropriate:
 - a. **Host:** The name of the upload host. If **Use secure connections (SSL)** is selected, the host name must match the host name in the certificate presented by the server. The host can be defined as an IPv4 or IPv6 address, or a domain name that resolves to an IPv4 or IPv6 address.
 - b. **Port:** If an IP address is entered for the host, specify a port number; the default is 80 for HTTP clients.

Note: For HTTPS, change the port to **443**.

- c. **Path:** The directory path where the access log facility is uploaded on the server.
 - d. **Username:** This is the username that is known on the host you are configuring.
 - e. **Change Password:** Change the password on the HTTP host; the Change Password dialog displays; enter and confirm the new password and click **OK**.
5. **Filename:** The **Filename** field is comprised of text and/or specifiers. The default filename includes specifiers and text that indicate the log name (%f), name of the external certificate used for encryption, if any (%c), the fourth parameter of the ProxySG IP address (%l), the date and time (Month: %m, Day: %d, Hour: %H, Minute: %M, Second: %S), and the .log or .gzip.log file extension.

Note: Be cautious if you change the **Filename** field. If an ongoing series of access log files are produced and you do not have time-specifiers in this field, each access log file produced overwrites the old file. Also, if you use more than one external certificate to encrypt logs, include the %c specifier in the **Filename** field to keep track of which external certificate can decrypt the uploaded log file.

6. **Local Time:** If you want the upload to reflect the local time it was uploaded instead of Universal Time Coordinate (UTC), select **Local Time**.
7. **Use secure connections (SSL):** Select this to create an HTTPS client. To create an HTTPS client, you must also create a key pair, import or create a certificate, and, if necessary, associate the key pair and certificate (called a keyring), with the SSL device profile.
8. Click **OK**.
9. Click **Apply**.

Editing the Custom client

To edit the custom client:

1. Select **Configuration > Access Logging > Logs > Upload Client**.
See "[Configuring the Upload Client](#)" on page 629 for configuration information.
2. Select **Custom client** from the **Client type** drop-down list. Click the **Settings** button.

Custom server connection:

3 Settings for: Primary Custom Server

4a Host: 10.1.2.3 Port: 69 4b

4c ☒ Use secure connections (SSL)

3. From the **Settings for** drop-down list, select to configure the primary or alternate custom server.
4. Fill in the server fields, as appropriate:
 - a. **Host:** Enter the IP address, in IPv4 format, of the upload destination. If **Use secure connections (SSL)** is selected, the host name must match the host name in the certificate presented by the server.

Note: Do not use a hostname instead of an IP address; doing so results in an error.

- b. **Port:** Specify a port number; the default is 69 for custom clients.
 - c. **Use secure connections (SSL):** Select this if you are using secure connections.
5. Click **OK**.
6. Click **Apply**.

Editing the Custom SurfControl Client

Use the Custom client to create an upload client that uploads information to SurfControl Reporter. Before you begin, verify that:

- ❑ You have created a log (see ["Creating and Editing an Access Log Facility"](#) on page 641).
- ❑ You have associated the SurfControl log format with the log you created (see ["Creating and Editing an Access Log Facility"](#) on page 641).

To edit the SurfControl client:

1. Select **Configuration > Access Logging > Logs > Upload Client**.
2. From the **Log** drop-down list, select the SurfControl log that you associated with the SurfControl log format.
3. Verify the **Save the log file as** radio button is set to **text file**, not **gzip file**.
4. Select **Custom client** from the **Client type** drop-down list.

Note: For specific information on managing upload clients, see ["Editing the Custom client"](#) on page 639.

5. Click the **Settings** button for that client.
6. Customize the upload client for SurfControl Reporter.
 - a. Enter the hostname, path, and username, if necessary, for the SurfControl Reporter server.
 - b. Ensure the filename extension is **.tmp** and not **.gzip** or **.log**. SurfControl only recognizes files with a **.tmp** extension.
 - c. If your SurfControl server supports SSL, select the **Use secure connections (SSL)** check box.
7. Click **OK**.
8. Click **Apply**.

Troubleshooting

- ❑ **Problem:** The ProxySG is uploading logs more frequently than expected.

Description: If access logging is enabled, logs can accrue on the ProxySG's hard drive even if the upload client is not configured for specific protocols (often the case if you configured streaming, IM, or P2P). Eventually the size of these combined logs, triggers the global **Start an Early upload** threshold (**Configuration > Access Logging > General > Global Settings**). The ProxySG attempts to upload all configured logs more often than expected. For example, a main log that is configured for upload every 24 hours starts to upload small portions of the main log every 10 minutes.

Solution: To prevent the access logs that do not have an upload client configured from triggering the **Start an Early upload** threshold, edit the default logs for each protocol that you do not need uploaded. Set them to **<None>** from the **Configuration > Access Logging > Logs > Upload Client** tab.

Chapter 28: Creating and Editing an Access Log Facility

This chapter describes how to modify existing log facilities for your needs. You can also create new log facilities for special circumstances, such as associating the SurfControl log format with a log facility.

Topics in this Chapter:

The following topics in this chapter include:

- ❑ ["Creating a Log Facility" on page 641](#)
- ❑ ["Editing an Existing Log Facility" on page 643](#)
- ❑ ["Deleting a Log Facility" on page 644](#)
- ❑ ["Disabling Access Logging for a Particular Protocol" on page 646](#)
- ❑ ["Configuring Global Settings" on page 646](#)


Creating a Log Facility

To create new log facilities, continue with the next section. To edit an existing log facility, skip to ["Configuring Global Settings" on page 646](#).

Note: Several log facilities have already been created. Before creating a new one, check the existing ones to see if they fit your needs. If you want to use a custom log format with the new log facility, you must create the log format before associating it with a log (see ["Creating Custom Access Log Formats" on page 649](#)).

To create a log facility:

1. Select **Configuration > Access Logging > Logs > Logs**.
2. The log facilities already created are displayed in the **Logs** tab. To create a new log, click **New**.



3. Fill in the fields as appropriate:

- a. **Log Name:** Enter a log facility name that is meaningful to you.

Note: The name can include specifiers from [Table 30–5](#) on page 664. For example, if you name the file:

- **AccLog**, the name will be **AccLog**
- **AccLog%C%m%d%H%M%S**, the name becomes **AccLog ProxySG_name month day hour min sec**
- **C%m%d**, the name becomes **ProxySG_name month day**
- **Y%m%d%C**, the name becomes **2008 month day ProxySG_name**

- b. **Log Format:** Select a log format from the drop-down list.

- c. **Description:** Enter a meaningful description of the log. It is used for display purposes only.

4. Fill in the **Log file limits** panel as appropriate. (You can edit these settings later. See "[Configuring Global Settings](#)" on page 646.)

- a. The maximum size for each remote log file (the file on the upload server) defaults to **0**, meaning that all data is sent to the same log file. If you set a maximum size, a new log file opens when the file reaches that size. This setting is valid for both periodic and continuous uploads.
- b. Specify a size that triggers an early upload—the maximum upload size varies depending on the size of the appliance disks (the maximum allowed upload threshold appears below this field).

5. Click **OK** to close the dialog.

6. Click **Apply**.

Editing an Existing Log Facility

Several facilities exist, each associated with a log format. For a description of the format, see ["Access Log Formats"](#) on page 657.

- ❑ **im** (Instant Messaging): Associated with the im format.
- ❑ **main**: Associated with the main format.
- ❑ **p2p** (Peer-to-Peer): Associated with the p2p format.
- ❑ **ssl**: Associated with the SSL format.
- ❑ **streaming**: Associated with the streaming format.

Use the following procedures to edit log facilities you have created.

Note: If you change the log format of a log, remember that ELFF formats require an ELFF header in the log (the list of fields being logged are mentioned in the header) and that non-ELFF formats do not require this header.

The format of data written to the log changes as soon as the format change is applied; for best practices, do a log upload before the format change and immediately after (to minimize the number of log lines in a file with mixed log formats).

Upload the log facility before you switch the format.

To edit an existing log facility:

1. Select **Configuration > Access Logging > Logs > General Settings**.

The screenshot displays the 'General Settings' configuration page for a log facility. The 'Log' dropdown is set to 'main'. The 'Log Settings' section includes a 'Log Format' dropdown set to 'main' and a 'Description' text field containing 'Use for protocols having no specific default log (such .'. The 'Log file limits' section shows two input fields: 'The maximum size of each remote file is' set to 0 megabytes, and 'Start an early upload if log reaches' set to 822 megabytes. A note at the bottom indicates 'Note: The maximum allowed upload threshold is: 1027 megabytes'. Callout boxes 2a, 2b, 2c, 3a, and 3b point to the 'Log', 'Log Format', 'Description', 'The maximum size of each remote file is', and 'Start an early upload if log reaches' fields respectively.

2. Fill in the fields as appropriate:
 - a. **Log:** Select an already-existing log facility from the **Log** drop-down list.
 - b. **Log Format:** Select the log format from the drop-down list.
 - c. **Description:** Enter a meaningful description of the log. (If you chose an existing log format, the default description for that log is displayed. You can change it.)


3. Fill in the **Log file limits** panel as appropriate:
 - a. The maximum size for each remote log file (the file on the upload server) defaults to **0**, meaning that all data is sent to the same log file. If you set a maximum size, a new log file opens when the file reaches that size. This setting is valid for both periodic and continuous uploads.
 - b. Specify a size that triggers an early upload—the maximum upload size varies depending on the size of the appliance disks (the maximum allowed upload threshold appears below this field).
4. Click **OK** to close the dialog.
5. Click **Apply**.

Deleting a Log Facility

You can delete a log facility through the Management Console.

To delete a log facility through the Management Console:

1. Select **Configuration > Access Logging > Logs**. All of the log facilities are displayed.



2. Select the log facility you want to delete and click **Delete**.
3. The Confirm Delete? dialog displays. Click **Ok**.

The log is successfully deleted when it is no longer displayed under **Logs**.

Associating a Log Facility with a Protocol

You can associate a log facility with a protocol at any point in the process. By default, new systems have specific protocols associated with specific logs. This allows you to begin access logging as soon as it is enabled.

Note: If you have a policy that defines protocol and log association, that policy overrides any settings you make here.

The following list shows the protocols supported and the default log facilities assigned to them, if any:

Table 28–1 Default Log Facility Assignments

Protocol	Assigned Default Log Facility
CIFS	cifs
Endpoint Mapper	main
Flash	streaming (for upgrades) bcreporterstreaming_v1 (for new systems)
FTP	main
HTTP	main
HTTPS-Reverse-Proxy	main (Set to the same log facility that HTTP is using upon upgrade.)
HTTPS-Forward-Proxy	ssl (If the facility for HTTP, TCP, or SOCKS is set before upgrade.)
Instant Messaging	im
MAPI	mapi
Peer to Peer	p2p
RealMedia/QuickTime	streaming (for upgrades) bcreporterstreaming_v1 (for new systems)
SOCKS	none
SSL	ssl (If the facility for HTTP, TCP or SOCKS is set before upgrade.)
TCP Tunnel	main
Telnet	main
Windows Media	streaming (for upgrades) bcreporterstreaming_v1 (for new systems)

Note: To disable access logging for a particular protocol, you must either disable the default logging policy for that protocol (see ["Disabling Access Logging for a Particular Protocol"](#) on page 646) or modify the access logging policy in VPM (refer to the *Visual Policy Manager Reference*).

To associate a log facility with a protocol:

1. Select **Configuration > Access Logging > General > Default Logging**.

2. Highlight the protocol you want to associate with a log facility and click **Edit**.
3. Select a log facility from the **Default Log** drop-down list.

Note: To disable access logging for that protocol, select **none**.

4. Click **OK** to close the dialog.
5. Click **Apply**.

Disabling Access Logging for a Particular Protocol

To disable access logging for a particular protocol:

1. Select **Configuration > Access Logging > General > Default Logging**.
2. Highlight the protocol to disable access logging and click **Edit**.
3. Select **none** from the drop-down menu.
4. Click **OK**.
5. Click **Apply**.

Configuring Global Settings

You might want to modify access log file sizes if, for example, the Management Console displays high disk usage for access logs (**Statistics > System > Resources > Disk Use**). To determine which access logs contribute to high disk usage, to look for a trend in log sizes, or for other troubleshooting scenarios, you can set the ProxySG appliance to take certain actions when the combined size of all access logs reaches specified global limits:

- ☐ Stop all access logging.
- ☐ Delete the oldest entries (overflow) from any of the log facilities regardless of which log caused the total size to reach the global limit.
- ☐ Attempt an early upload of the log that caused the total size to reach the global limit.

You can also upload all configured access logs immediately.

After monitoring which logs the appliance uploads over a period of time, you can modify the amount of space allocated to individual logs for early upload.

To specify global settings:

1. Select **Configuration > Access Logging > General > Global Settings**.

2. Enter global limits in the **Global Log File Limits** section.

The screenshot shows the 'Global Settings' window with the 'Global Log File Limits' section. The 'Limit total system access logging to' field is set to 1028 megabytes. The 'stop logging' radio button is selected. The 'Start an early upload if total logging reaches' field is set to 822 megabytes. The 'Global Upload' section has 'Upload All' and 'Cancel All' buttons. The 'Apply' button is at the bottom.

- a. Enter the maximum total size of all log files. This is the sum of the sizes of all the individual logs.
- b. Specify what the appliance should do when the total log size reaches the maximum:
 - Stop all access logging and attempt an immediate upload.
 - Delete the oldest entries (overflow) from any of the log facilities regardless of which log caused the total size to reach the global limit.

Note: To ensure that the appliance uploads older log entries before they are deleted, make sure that you have correctly specified an upload server, an early upload threshold for each individual log, and a global early upload threshold. Inability to connect to the upload server (for example, due to incorrect settings or network issues) will result in data loss.

- c. Specify the sum total of all log sizes to trigger an early upload. The appliance attempts to upload the log that caused the total log size to reach the global limit.

Individual logs each have their own early upload setting, which remain in effect even if you specify a global value. This means that, provided that upload is configured and working, the global early upload threshold could trigger a log upload before the file size reaches the threshold defined in the specific log facility.

3. The **Global Upload** section allows you to perform actions on all available log facilities.
 - Click **Upload All** to upload all logs immediately, and click **OK** to confirm the action.
 - Click **Cancel** to prevent further attempts to upload all log facilities and cancel all uploads that are in progress.
4. Click **Apply**.

Chapter 29: Creating Custom Access Log Formats

This chapter describes the default access log formats and describes how to create customized access log formats.

Topics in this Chapter:

This chapter includes information about the following topics:

- ❑ ["Default Access Log Formats"](#) on page 649
- ❑ ["Creating a Custom or ELFF Log Format"](#) on page 652

Default Access Log Formats

Several log formats ship with the SGOS software, and they might be sufficient for your needs. If the formats that exist do not meet your needs, you can create a custom or ELFF format and specify the string and other qualifiers used, as described in ["Creating a Custom or ELFF Log Format"](#) on page 652.

Note: Reserved log formats cannot be edited or modified in any way. If you wish to create a custom log format based on an existing reserved log format, see ["Creating a Custom or ELFF Log Format"](#) on page 652.

For a description of each value in the log, see ["Access Log Formats"](#) on page 657.

- ❑ **cifs:** This is an ELFF format with the custom strings of

```
date time c-ip c-port r-ip r-port s-action s-ip cs-auth-group cs-  
username x-client-connection-bytes x-server-connection-bytes x-  
server-adn-connection-bytes x-cifs-method x-cifs-client-read-  
operations x-cifs-client-write-operations x-cifs-client-other-  
operations x-cifs-server-operations x-cifs-error-code x-cifs-server  
x-cifs-share x-cifs-path x-cifs-orig-path x-cifs-client-bytes-read  
x-cifs-server-bytes-read x-cifs-bytes-written x-cifs-uid x-cifs-tid  
x-cifs-fid x-cifs-file-size x-cifs-file-type
```

- ❑ **mapi:** This is an ELFF format with the custom strings of

```
date time c-ip c-port r-ip r-port x-mapi-user x-mapi-method cs-bytes  
sr-bytes rs-bytes sc-bytes x-mapi-cs-rpc-count x-mapi-sr-rpc-count  
x-mapi-rs-rpc-count x-mapi-sc-rpc-count s-action cs-username cs-  
auth-group s-ip
```

- ❑ **im (Instant Messaging):** This is an ELFF format with the custom strings of:

```
date time c-ip cs-username cs-auth-group cs-protocol x-im-method x-  
im-user-id x-im-user-name x-im-user-state x-im-client-info x-im-  
buddy-id x-im-buddy-name x-im-buddy-state x-im-chat-room-id x-im-  
chat-room-type x-im-chat-room-members x-im-message-text x-im-  
message-size x-im-message-route x-im-message-type x-im-file-path x-  
im-file-size s-action
```

- ❑ **main:** This is an ELFF format with custom strings of:

```
date time time-taken c-ip sc-status s-action sc-bytes cs-bytes cs-  
method cs-uri-scheme cs-host cs-uri-port cs-uri-path cs-uri-query cs-  
username cs-auth-group s-supplier-name rs(Content-Type) cs(Referer)  
cs(User-Agent) sc-filter-result cs-categories x-virus-id s-ip
```

- ❑ **nrsa:** This is a reserved format that cannot be edited. The NCSA/Common format contains the following strings:

```
remotehost rfc931 authuser [date] "request" status bytes
```

The ELFF/custom access log format strings that represent the strings above are:

```
$(c-ip) - $(cs-username) $(localtime) $(cs-request-line) $(sc-status)  
$(sc-bytes)
```

- ❑ **p2p:** This is an ELFF format with custom strings of:

```
date time c-ip c-dns cs-username cs-auth-group cs-protocol x-p2p-  
client-type x-p2p-client-info x-p2p-client-bytes x-p2p-peer-bytes  
duration s-action
```

- ❑ **squid:** This is a reserved format that cannot be edited. You can create a new SQUID log format using custom strings. The default SQUID format is SQUID-1.1 and SQUID-2 compatible.

SQUID uses several definitions for its field formats:

```
SQUID-1:time elapsed remotehost code/status/peerstatus bytes method  
URL
```

```
SQUID-1.1: time elapsed remotehost code/status bytes method URL rfc931  
peerstatus/peerhost type
```

SQUID-2 has the same fields as SQUID-1.1, although some of the field values have changed.

- ❑ **ssl:** This is an ELFF format with custom strings of:

```
date time time-taken c-ip s-action x-rs-certificate-validate-status x-  
rs-certificate-observed-errors x-cs-ocsp-error x-rs-ocsp-error cs-host  
s-supplier-name x-rs-connection-negotiated-ssl-version x-rs-  
connection-negotiated-cipher x-rs-connection-negotiated-cipher-size x-  
rs-certificate-hostname x-rs-certificate-hostname-category x-cs-  
connection-negotiated-ssl-version x-cs-connection-negotiated-cipher x-  
cs-connection-negotiated-cipher-size x-cs-certificate-subject s-ip s-  
sitename
```

- ❑ **streaming:** This is an ELFF format with custom strings of:

```
c-ip date time c-dns cs-uri-scheme cs-host cs-uri-port cs-uri-path cs-  
uri-query c-starttime x-duration c-rate c-status c-playerid c-  
playerversion c-playerlanguage cs(User-Agent) cs(Referer) c-hostexe c-  
hostexeever c-os c-osversion c-cpu filelength filesize avgbandwidth  
protocol transport audiocodec videocodec channelURL sc-bytes c-bytes  
s-pkts-sent c-pkts-received c-pkts-lost-client c-pkts-lost-net c-pkts-  
lost-cont-net c-resendreqs c-pkts-recovered-ECC c-pkts-recovered-  
resent c-buffercount c-totalbuffertime c-quality s-ip s-dns s-  
totalclients s-cpu-util x-cache-user s-session-id x-cache-info x-  
client-address s-action
```

- ❑ **bcreportercifs_v1** is designed to for Proxy deployments that use ADN to transfer data with CIFS, and send that access information to Symantec Reporter. This is a reserved format and cannot be edited. The included access logging fields are:

```
date time c-ip c-port r-ip r-port s-action s-ip cs-auth-group cs-username x-client-connection-bytes x-server-connection-bytes x-server-adn-connection-bytes x-cifs-method x-cifs-client-read-operations x-cifs-client-write-operations x-cifs-client-other-operations x-cifs-server-operations x-cifs-error-code x-cifs-server x-cifs-share x-cifs-path x-cifs-orig-path x-cifs-client-bytes-read x-cifs-server-bytes-read x-cifs-bytes-written x-cifs-uid x-cifs-tid x-cifs-fid x-cifs-file-size x-cifs-file-type
```

- ❑ **bcreportermain_v1** is designed to send HTTP access information to Symantec Reporter. This is a reserved format and cannot be edited. The included access logging fields are:

```
date time time-taken c-ip cs-username cs-auth-group x-exception-id sc-filter-result cs-categories cs(Referer) sc-status s-action cs-method rs(Content-Type) cs-uri-scheme cs-host cs-uri-port cs-uri-path cs-uri-query cs-uri-extension cs(User-Agent) s-ip sc-bytes cs-bytes x-virus-id x-bluecoat-application-name x-bluecoat-application-operation
```

Note: In version 6.5.9.2 and later, **bcreportermain_v1** includes the `x-bluecoat-transaction-uuid`, `x-icap-reqmod-header (X-ICAP-Metadata)`, and `x-icap-respmo-header (X-ICAP-Metadata)` fields.

- ❑ **bcreporterssl_v1** is designed to send HTTPs access information to Symantec Reporter. This is a reserved format and cannot be edited. The included access logging fields are::

```
date time time-taken c-ip cs-username cs-auth-group x-exception-id sc-filter-result cs-categories sc-status s-action cs-method rs(Content-Type) cs-uri-scheme cs-host cs-uri-port cs-uri-extension cs(User-Agent) s-ip sc-bytes cs-bytes x-virus-id x-rs-certificate-observed-errors x-cs-ocsp-error x-rs-ocsp-error x-rs-connection-negotiated-cipher-strength x-rs-certificate-hostname x-rs-certificate-hostname-category
```

- ❑ **bcreporterstreaming_v1** is designed to send streaming media access information to Symantec Reporter. This is a reserved format and cannot be edited. The included access logging fields are::

```
date time time-taken c-ip sc-status s-action sc-bytes rs-bytes cs-method cs-uri-scheme cs-host cs-uri-port cs-uri-path cs-uri-query cs-username cs-auth-group cs(Referer) cs(User-Agent) c-starttime filelength filesize avgbandwidth x-rs-streaming-content x-streaming-rtmp-app-name x-streaming-rtmp-stream-name x-streaming-rtmp-swf-url x-streaming-rtmp-page-url s-ip s-dns s-session-id x-cache-info
```

- ❑ **bcreporterwarp_v1** is designed to send reverse proxy access information to Symantec Reporter. This is a reserved format and cannot be edited. The included access logging fields are:

```

date time time-taken c-ip cs-username cs-auth-group x-exception-id
cs(Referer) sc-status s-action cs-method rs(Content-Type) cs-uri-
scheme cs-host cs-uri-port cs-uri-path cs-uri-query cs-uri-extension
cs(User-Agent) s-ip sc-bytes cs-bytes x-virus-id x-cs-client-ip-
country x-risk-category x-risk-score x-user-x509-serial-number x-user-
x509-subject rs-bytes x-cs-client-effective-ip x-cs-client-effective-
ip-country cs(X-Forwarded-For) rs-service-time-taken r-ip

```

Creating a Custom or ELFF Log Format


First, decide what protocols and log formats to use, and determine the logging policy and the upload schedule. Then perform the following:

- ❑ Associate a log format with the log facility.
- ❑ Associate a log facility with a protocol and/or create policies for protocol association and to manage the access logs and generate entries in them (if you do both, policy takes precedence).
- ❑ Determine the upload parameters for the log facility.


For more information, see ["Default Access Log Formats"](#) on page 649.

To create or edit the log format:

1. Select **Configuration > Access Logging > Formats**.



2. Click **New** (or highlight a format and click **Edit**). The Create Format dialog displays. If you select an unconfigurable format, you receive an error message.



3. Create or modify the format:

- a. Give the format a meaningful name.
- b. Select **Custom format string** (to manually add your own format field) or **W3C ELFF** (to customize using the standard format fields).
- c. Add log formats or remove from the current list.

Note: ELFF strings cannot start with spaces.

The access log ignores any ELFF or custom format fields it does not understand. In a downgrade, the format still contains all the fields used in the upgraded version, but only the valid fields for the downgraded version display any information.


- d. Click **Test Format** to test whether the format-string syntax is correct. A line displays below the field that indicates that testing is in progress and then gives a result, such as **Format is valid**.
 - e. From the **Multiple-valued header policy** drop-down list, select a header to log: **Log last header**, **log first header**, **log all headers**. This allows you to determine what happens with HTTP-headers that have multiple headers.
 - f. Click **OK**.
4. Click **Apply**.

Creating Custom Log Formats Based on Reserved Log Formats

There might be instances where the reserved log format is insufficient for your purposes and requires either a log format extension or reduction. Although the reserved log formats cannot be directly manipulated, you can create new custom log formats based on these reserved log formats.

To copy a reserved log format into a custom schema:

1. Select an existing reserved log format that contains the format string you wish to copy.



2. Click **Edit/View**. The **View Format** dialog box appears.



3. Highlight the portion of the string that you wish to copy and use a keyboard shortcut to copy the text onto the clipboard.

Note: Be aware that you cannot copy and paste selections using the right mouse button from within the Management Console; you must use keyboard shortcuts.

4. Click **Cancel** to close the window.
5. Click **New** (or highlight an existing format and click **Edit**). The **Create Format** (or **Edit Format**, if you are editing an existing format) dialog displays.

Create Format

Format Settings:

Format Name:

☐ Custom format string (specify below)

☒ W3C Extended Log File Format (ELFF) string (specify below)

Multiple-valued header policy:

6. Select the format string field (if there is an existing string, place the cursor where you want to insert the string) and paste the string from the clipboard using a keyboard shortcut.
7. Continue from step 3 from ["To create or edit the log format:"](#) on page 652.

Related CLI Syntax to Manage Access Logging

Some options for custom access log formats cannot be configured in the Management Console. Refer to the following commands to manage access logging.

- To enter configuration mode:

```
SGOS#(config) access-log
```

The following subcommands are available:

```
SGOS#(config access-log) create log log_name
SGOS#(config access-log) create format format_name
SGOS#(config access-log) cancel-upload all
SGOS#(config access-log) cancel-upload log log_name
SGOS#(config access-log) default-logging {cifs | epmapper | ftp | http
| https-forward-proxy | https-reverse-proxy | im | mapi | mms | p2p |
rtsp | socks | ssl | tcp-tunnel | telnet} log_name
SGOS#(config access-log) delete log log_name
SGOS#(config access-log) delete format format_name
SGOS#(config access-log) disable
SGOS#(config access-log) early-upload megabytes
SGOS#(config access-log) edit log log_name—changes the prompt to
SGOS#(config edit log log_name)
SGOS#(config access-log) edit format format_name—changes the prompt to
SGOS#(config edit format format_name)
SGOS#(config access-log) enable
SGOS#(config access-log) exit
SGOS#(config access-log) max-log-size megabytes
SGOS#(config access-log) no default-logging {cifs | epmapper | ftp |
http | https-forward-proxy | https-reverse-proxy | im | mapi | mms |
p2p | rtsp | socks | ssl | tcp-tunnel | telnet}
SGOS#(config access-log) overflow-policy delete
SGOS#(config access-log) overflow-policy stop
SGOS#(config access-log) upload all
SGOS#(config access-log) upload log log_name
SGOS#(config access-log) view
SGOS#(config access-log) view [log [brief | log_name]]
SGOS#(config access-log) view [format [brief | format_name]]
SGOS#(config access-log) view [statistics [log_name]]
SGOS#(config access-log) view [default-logging]
```


Chapter 30: Access Log Formats

This chapter describes the access log formats that are created by ProxySG:

- ❑ "Custom or W3C ELFF Format"
- ❑ "SQUID-Compatible Format" on page 661
- ❑ "NCSA Common Access Log Format" on page 664

ELFF is a log format defined by the W3C that contains information about Windows Media and RealProxy logs.

The ProxySG can create access logs with any one of six formats. Four of the six are reserved formats and cannot be configured. However, you can create additional logs using custom or ELFF format strings.

When using an ELFF or custom format, a blank field is represented by a dash character. When using the SQUID or NCSA log format, a blank field is represented according to the standard of the format.

Custom or W3C ELFF Format

The W3C Extended Log File Format (ELFF) is a subset of the Blue Coat Systems format. The ELFF format is specified as a series of space delimited fields. Each field is described using a text string. The types of fields are described in the following table.

Table 30–1 Field Types

Field Type	Description
Identifier	A type unrelated to a specific party, such as date and time.
prefix-identifier	Describes information related to a party or a transfer, such as <code>c-ip</code> (client's IP) or <code>sc-bytes</code> (how many bytes were sent from the server to the client)
prefix (header)	Describes a header data field. The valid prefixes are: <code>c</code> = Client <code>s</code> = Server <code>r</code> = Remote <code>sr</code> = Server to Remote <code>cs</code> = Client to Server <code>sc</code> = Server to Client <code>rs</code> = Remote to Server

ELFF formats are created by selecting a corresponding custom log format using the table below. Unlike the Blue Coat custom format, ELFF does not support character strings and require a space between fields.

Selecting the ELFF format does the following:

- ❑ Puts one or more W3C headers into the log file. Each header contains the following lines:


```
#Software: SGOS x.x.x
#Version: 1.0
#Date: 2002-06-06 12:12:34
#Fields: date time cs-ip...
```
- ❑ Changes all spaces within fields to + or %20. The ELFF standard requires that spaces only be present between fields.

ELFF formats are described in the following table.

Table 30–2 Blue Coat Custom Format and Extended Log File Format

Blue Coat Custom Format	Extended Log File Format	Description
space character	N/A	Multiple consecutive spaces are compressed to a single space.
%	-	Denotes an expansion field.
%%	-	Denotes '%' character.
%a	c-ip	IP address of the client
%b	sc-bytes	Number of bytes sent from appliance to client
%c	rs (Content-Type)	Response header: Content-Type
%d	s-supplier-name	Hostname of the upstream host (not available for a cache hit)
%e	time-taken	Time taken (in milliseconds) to process the request (from the first byte of client request data received by the proxy, to the last byte sent by the proxy to the client, including all of the delays by ICAP, and so on).
%f	sc-filter-category	Content filtering category of the request URL
%g	timestamp	Unix type timestamp
%h	c-dns	Hostname of the client (uses the client's IP address to avoid reverse DNS)
%i	cs-uri	The 'log' URL.
%j	-	[Not used.]
%k	-	[Not used.]
%l	x-bluecoat-special-empty	Resolves to an empty string
%m	cs-method	Request method used from client to appliance

Table 30–2 Blue Coat Custom Format and Extended Log File Format (Continued)

Blue Coat Custom Format	Extended Log File Format	Description
%n	-	[Not used.]
%o	-	[Not used.]
%p	r-port	Port from the outbound server URL
%q	-	[Not used.]
%r	cs-request-line	First line of the client's request
%s	sc-status	Protocol status code from appliance to client
%t	gmttime	GMT date and time of the user request in format: [DD/MM/YYYY:hh:mm:ss GMT]
%u	cs-user	Qualified username for NTLM. Relative username for other protocols
%v	cs-host	Hostname from the client's request URL. If URL rewrite policies are used, this field's value is derived from the 'log' URL
%w	s-action	What type of action did the ProxySG take to process this request (see " Action Field Values " on page 661)
%x	date	GMT Date in YYYY-MM-DD format
%y	time	GMT time in HH:MM:SS format
%z	s-icap-status	ICAP response status
%A	cs (User-Agent)	Request header: User-Agent
%B	cs-bytes	Number of bytes sent from client to appliance
%C	cs (Cookie)	Request header: Cookie
%D	s-supplier-ip	IP address used to contact the upstream host (not available for a cache hit)
%E	-	[Not used.]
%F	-	[Not used.]
%G	-	[Not used.]
%H	s-hierarchy	How and where the object was retrieved in the cache hierarchy.
%I	s-ip	IP address of the appliance on which the client established its connection
%J	-	[Not used.]

Table 30–2 Blue Coat Custom Format and Extended Log File Format (Continued)

Blue Coat Custom Format	Extended Log File Format	Description
%K	-	[Not used.]
%L	localtime	Local date and time of the user request in format: [DD/MMM/YYYY:hh:mm:ss +nnnn]
%M	-	[Not used.]
%N	s-computername	Configured name of the appliance
%O	-	[Not used.]
%P	s-port	Port of the appliance on which the client established its connection
%Q	cs-uri-query	Query from the 'log' URL.
%R	cs (Referer)	Request header: Referer
%S	s-sitename	The service type used to process the transaction
%T	duration	Time taken (in seconds) to process the request
%U	cs-uri-path	Path from the 'log' URL. Does not include query.
%V	cs-version	Protocol and version from the client's request, e.g. HTTP/1.1
%W	sc-filter-result	Content filtering result: Denied, Proxied or Observed
%X	cs (X-Forwarded-For)	Request header: X-Forwarded-For
%Y	-	[Not used.]
%Z	s-icap-info	ICAP response information

Example Access Log Formats

```

Squid log format: %g %e %a %w/%s %b %m %i %u %H/%d %c
NCSA common log format: %h %l %u %t "%r" %s %b
NCSA extended log format: %h %l %u %L "%r" %s %b "%R" "%A"
Microsoft IIS format: %a, -, %x, %y, %S, %N, %I, %e, %b, %B, %s, 0, %m, %U, -

```

The Blue Coat custom format allows any combination of characters and format fields. Multiple spaces are compressed to a single space in the actual access log. You can also enter a string, such as `My default is %d`. The ProxySG goes through such strings and finds the relevant information. In this case, that information is `%d`.

SQUID-Compatible Format

The SQUID-compatible format contains one line for each request. For SQUID-1.1, the format is:

```
time elapsed remotehost code/status bytes method URL rfc931
peerstatus/peerhost type
```

For SQUID-2, the columns stay the same, though the content within might change a little.

Action Field Values

Table 1–3 describes the possible values for the `s-action` field.

Table 30–3 Action Field Values

Value	Description
ACCELERATED	(SOCKS only) The request was handed to the appropriate protocol agent for handling.
ALLOWED	An FTP method (other than the data transfer method) is successful.
DENIED	Policy denies a method. A <code>DENIED</code> <code>s-action</code> value is returned for CIFS, Endpoint Mapper, MAPI, FTP, IM, P2P, Shell Proxy, SOCKS proxy, streaming proxies, and SSL proxy when policy denies a request. When the same kind of denial happens in the HTTP proxy, <code>TCP_DENIED</code> is reported.
FAILED	An error or failure occurred.
LICENSE_EXPIRED	(SOCKS only) The request could not be handled because the associated license has expired.
TUNNELED	Successful data transfer operation.
TCP_	Refers to requests on the HTTP port.
TCP_ACCELERATED	For <code>CONNECT</code> tunnels that are handed off to the following proxies: HTTP, SSL, Endpoint mapper, and P2P for BitTorrent/EDonkey/Gnutella.
TCP_AUTH_HIT	The requested object requires upstream authentication, and was served from the cache.
TCP_AUTH_HIT_RST	The requested object requires upstream authentication, but the client connection was reset before the complete response was delivered.
TCP_AUTH_MISS	The requested object requires upstream authentication, and was not served from the cache. This is part of CAD (Cached Authenticated Data).
TCP_AUTH_MISS_RST	The requested object requires upstream authentication, and was not served from the cache; the client connection was reset before the complete response was delivered.

Table 30–3 Action Field Values (Continued)

Value	Description
TCP_AUTH_FORM	Forms-based authentication is being used and a form challenging the user for credentials is served in place of the requested content. Note: Upon submission of the form, another access log entry is generated to indicate the status of the initial request.
TCP_AUTH_REDIRECT	The client was redirected to another URL for authentication.
TCP_BYPASSED	A TCP-Tunnel connection was bypassed because an upstream ADN concentrator was not discovered; this can occur only when the bypass-if-no-concentrator feature is enabled and all conditions for activating the feature are met. See "Discovery of Upstream Concentrators" on page 718.
TCP_CLIENT_REFRESH	The client forces a revalidation with the origin server with a Pragma: no-cache. If the server returns 304 Not Modified, this appears in the Statistics:Efficiency file as In Cache, verified Fresh.
TCP_CLIENT_REFRESH_RST	The client forces a revalidation with the origin server, but the client connection was reset before the complete response was delivered.
TCP_DENIED	Access to the requested object was denied by a filter.
TCP_ERR_MISS	An error occurred while retrieving the object from the origin server.
TCP_HIT	A valid copy of the requested object was in the cache.
TCP_HIT_RST	A valid copy of the requested object was in the cache, but the client connection was reset before the complete response was delivered.
TCP_LOOP	The current connection is dropped because the upstream connection would result in a looped connection.
TCP_MEM_HIT	The requested object was, in its entirety, in RAM.
TCP_MISS	The requested object was not in the cache.
TCP_MISS_RST	The requested object was not in the cache; the client connection was reset before the complete response was delivered.
TCP_NC_MISS	The object returned from the origin server was non-cacheable.

Table 30–3 Action Field Values (Continued)

Value	Description
TCP_NC_MISS_RST	The object returned from the origin server was non-cacheable; the client connection was reset before the complete response was delivered.
TCP_PARTIAL_MISS	The object is in the cache, but retrieval from the origin server is in progress.
TCP_PARTIAL_MISS_RST	The object is in the cache, but retrieval from the origin server is in progress; the client connection was reset before the complete response was delivered.
TCP_POLICY_REDIRECT	The client was redirected to another URL due to policy.
TCP_REFRESH_HIT	A GIMS request to the server was forced and the response was 304 Not Modified; this appears in the Statistics:Efficiency file as In Cache, verified Fresh.
TCP_REFRESH_HIT_RST	A GIMS request to the server was forced and the response was 304 Not Modified; the client connection was reset before the complete response was delivered.
TCP_REFRESH_MISS	A GIMS request to the server was forced and new content was returned.
TCP_REFRESH_MISS_RST	A GIMS request to the server was forced and new content was returned, but the client connection was reset before the complete response was delivered.
TCP_RESCAN_HIT	The requested object was found in the cache but was rescanned because the virus-scanner-tag-id in the object was different from the current scanner tag.
TCP_RESCAN_HIT_RST	The requested object was rescanned (see TCP_RESCAN_HIT) but the client connection was reset before the complete response was delivered.
TCP_SPLASHED	The user was redirected to a splash page.
TCP_SWAPFAIL	The object was believed to be in the cache, but could not be accessed.
TCP_TUNNELED	The CONNECT method was used to tunnel this request (generally proxied HTTPS).
TCP_WEBSOCKET	(Introduced in SGOS 6.5.5.7) The request was a WebSocket upgrade request. You can determine if the traffic was plain WebSocket or secure WebSocket by looking at the scheme (HTTP or HTTPS).

NCSA Common Access Log Format

The common log format contains one line for each request. The format of each log entry is shown below:

```
remotehost rfc931 authuser [date] "request" status bytes
```

Each field is described in the following table.

Table 30–4 Log Entry Fields

Field Name	Description
remotehost	DNS hostname or IP address of remote server.
rfc931	The remote log name of the user. This field is always —.
authuser	The username as which the user has authenticated himself.
[date]	Date and time of the request.
"request"	The request line exactly as it came from the client.
status	The HTTP status code returned to the client.
bytes	The content length of the document transferred.

Access Log Filename Formats

The following table details the specifiers for the access log upload filenames.

Table 30–5 Specifiers for Access Log Upload Filenames

Specifier	Description
%%	Percent sign.
%a	Abbreviated weekday name.
%A	Full weekday name.
%b	Abbreviated month name.
%B	Full month name.
%c	The certificate name used for encrypting the log file (expands to nothing in non-encrypted case).
%C	The ProxySG name.
%d	Day of month as decimal number (01 - 31).
%f	The log name.
%H	Hour in 24-hour format (00 - 23).
%i	First IP address of the ProxySG, displayed in x_x_x_x format, with leading zeros removed.
%I	Hour in 12-hour format (01 - 12).
%j	Day of year as decimal number (001 - 366).

Table 30–5 Specifiers for Access Log Upload Filenames (Continued)

%l	The fourth (last) octet in the ProxySG IP address (For example, for the IP address 10.11.12.13, %l would be 13)
%m	Month as decimal number (01 – 12).
%M	Minute as decimal number (00 – 59).
%p	Current locale's A.M./P.M. indicator for 12-hour clock.
%S	Second as decimal number (00 – 59).
%U	Week of year as decimal number, with Sunday as first day of week (00 – 53).
%v	Milliseconds; usually sed in conjunction with %H%M%S to get more accuracy in the log filename. <i>Available in SGOS 6.5.4.1 and higher</i>
%w	Weekday as decimal number (0 – 6; Sunday is 0).
%W	Week of year as decimal number, with Monday as first day of week (00 – 53).
%y	Year without century, as decimal number (00 – 99).
%Y	Year with century, as decimal number.
%z, %Z	Time-zone name or abbreviation; no characters if time zone is unknown.

Fields Available for Creating Access Log Formats

Refer to the *ProxySG Log Fields and CPL Substitutions Reference*:

<https://www.symantec.com/docs/DOC11251>

Chapter 31: Statistics

This chapter describes the statistics displayed in the Management Console. Statistics present a graphical view of the status for many system operations. This chapter also refers to NetFlow, which is available in the CLI only.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ ["Viewing the Traffic Mix Report"](#) on page 667
- ❑ ["Viewing NetFlow Statistics"](#) on page 672
- ❑ ["Viewing Traffic History"](#) on page 673
- ❑ ["Supported Proxies and Services"](#) on page 675
- ❑ ["Viewing the Application Mix Report"](#) on page 676
- ❑ ["Viewing the Application History Report"](#) on page 681
- ❑ ["Viewing System Statistics"](#) on page 682
- ❑ ["Active Sessions—Viewing Per-Connection Statistics"](#) on page 690

Viewing the Traffic Mix Report

The Traffic Mix report allows you to view traffic distribution and bandwidth statistics for traffic running through the ProxySG. You can break down the data according to proxy type or service name across various time periods.

The report has three parts to it:

- ❑ Line graph showing bandwidth usage or gain (see ["Viewing Bandwidth Details for Proxies or Services"](#) on page 668)
- ❑ Pie graph showing traffic distribution of proxies or services (see ["Viewing Traffic Distribution"](#) on page 670)
- ❑ Statistical table listing client/server bytes and savings for each proxy / service (see ["Viewing Per-Proxy or Per-Service Statistics"](#) on page 671)




Figure 31–1 Traffic Mix Report

Viewing Bandwidth Details for Proxies or Services

To see how much bandwidth is attributed to various proxies or services used on your network over the last hour, day, week, month, or year, view the line graph on the Traffic Mix report. This graph also allows you to analyze how much bandwidth you are gaining from optimization of proxy or service traffic.

To view bandwidth statistics for proxies or services:

1. Select **Statistics > Traffic Details > Traffic Mix**.
2. Select either **Service** or **Proxy**.

3. (Optional) Clear the **Include bypassed bytes** check box if you don't want to include bypassed traffic in the graphs, statistics, and calculations; this would allow you to get a clearer view of traffic that is intercepted.

4. To see the bandwidth rate of service/proxy traffic, select the **BW Usage** tab (underneath the line graph).

The green area represents client data, the blue area is server data, and the brown is bypassed bytes (if included).

5. To see how much bandwidth is gained due to optimization of server/proxy traffic, select the **BW Gain** tab.

The line graph indicates the bandwidth gain due to optimizations, averaged over the time interval, expressed as a multiple (for example, 2x means that twice the amount of bandwidth is available).

6. Select the time period you are interested in from the **Duration** drop-down list.

The graphs and statistics automatically update to reflect the time period you selected. Thereafter, the chart data automatically updates every 60 seconds.

Hover the mouse cursor over the chart data to view detailed values.




Figure 31-2 Traffic Mix Statistics— displayed when the cursor hovers over chart data

The values that display when you hover the mouse cursor over the chart data can include:

- ❑ **C** = Client-side traffic data rate. This statistic represents the data rate calculated (to and from the client) on the client-side connection. Data rate is represented by units of bits per second (bps) from measurements that are sampled at one-minute intervals. All application protocol-level bytes are counted, including application-protocol overhead such as HTTP and CIFS headers.
- ❑ **S** = Server-side traffic data rate. This statistic represents the data rate calculated (to and from the server) on the server-side connection. The data rate is represented by units of bits per second (bps) from measurements that are sampled at one-minute intervals. All application-level bytes are counted, including application overhead such as HTTP and CIFS headers.
- ❑ **Unopt** = Unoptimized traffic data rate. This statistic reflects the data rate of original traffic served to/from the client or server prior to or subsequent to ADN optimization. The data rate is represented by units of bits per second (bps).

- ❑ **Opt** = Optimized traffic data rate. This statistic reflects the data rate of ADN-optimized traffic. Data rate is represented by units of bits per second (bps).
- ❑ **B** = Bypassed traffic data rate. This statistic reflects that data rate of bypassed traffic (traffic that is not intercepted by ProxySG services). The data rate is represented by units of bits per second (bps).
- ❑ **Gain** = Bandwidth Gain. This statistic, representing the overall bandwidth benefit achieved by object and byte caching, compression, protocol optimization, and object caching, is computed by the ratio:

`client bytes / server bytes`

and represented as a unit-less multiplication factor. Bandwidth-gain values are computed at one-minute intervals to correspond to the one-minute sampling of client and server bytes. For example, if server bytes displayed as 10kbps and client bytes was 90kbps, the bandwidth gain is represented as 9x.

- ❑ **Savings** = Bandwidth Savings. This statistic, representing the overall bandwidth savings achieved over the WAN by utilizing object and byte caching, protocol optimization, and compression, is computed by

`(client bytes - server bytes) / client bytes`

and presented as a percentage. The Savings value provides a relative percentage of bandwidth savings on the WAN link, with 100% indicating no WAN traffic at all (no server bytes) and 0% indicating that no savings were achieved by client bytes equaling server bytes. Utilizing the numbers from the above example, the equivalent savings would be $8/9 = 0.89 = 89\%$.

See Also

- ❑ ["Viewing Traffic Distribution"](#)
- ❑ ["Viewing Per-Proxy or Per-Service Statistics"](#)
- ❑ ["Clearing the Statistics"](#)
- ❑ ["About Bypassed Bytes"](#)
- ❑ ["About the Default Service Statistics"](#)

Viewing Traffic Distribution

The pie chart on the Traffic Mix report shows the distribution of service/proxy traffic over the last hour, day, week, month, or year. You can look at either client bytes or server bytes.

To view a pie chart showing distribution of service/proxy traffic:

1. Select **Statistics > Traffic Details > Traffic Mix**.
2. Select **Client Bytes** or **Server Bytes** (tabs underneath the pie chart).
3. Select a time period from the **Duration** drop-down list.

The pie chart displays data for the seven services/proxies with the most traffic during the selected time period; all other service/proxy statistics are placed into the **Other** category.

For a list of supported proxies and services, see ["Supported Proxies and Services"](#) on page 675.

See Also

- ❑ ["Viewing Bandwidth Details for Proxies or Services"](#)
- ❑ ["Viewing Per-Proxy or Per-Service Statistics"](#)
- ❑ ["Clearing the Statistics"](#)
- ❑ ["About Bypassed Bytes"](#)
- ❑ ["About the Default Service Statistics"](#)

Viewing Per-Proxy or Per-Service Statistics

The table of statistics at the bottom of the Traffic Mix report lists the following details for each proxy/service during the selected time period:

- ❑ **Client Bytes**—The data rate calculated (to and from the client) on the client-side connection, measured in bits per second (bps)
- ❑ **Server Bytes**—The data rate calculated (to and from the server) on the server-side connection, measured in bps
- ❑ **Bypassed Bytes**—The data rate of bypassed traffic (traffic that is not intercepted by ProxySG services), measured in bps
- ❑ **Savings**—Bandwidth savings achieved over the WAN by utilizing object and byte caching, protocol optimization, and compression; presented as a percentage. The formula is:

$$(\text{client bytes} - \text{server bytes}) / \text{client bytes}$$
See Also

- ❑ ["Viewing Bandwidth Details for Proxies or Services"](#)
- ❑ ["Viewing Traffic Distribution"](#)
- ❑ ["Clearing the Statistics"](#)
- ❑ ["About Bypassed Bytes"](#)
- ❑ ["About the Default Service Statistics"](#)

Clearing the Statistics

To reset traffic mix statistics, select **Maintenance > System and Disks > Tasks**, and click **Clear the trend statistics**.

About Bypassed Bytes

Bypassed bytes are bytes that are not intercepted by a service or proxy. By default, bypassed bytes are included in the traffic mix views. When evaluating traffic statistics for potential optimization, it can be useful to include or exclude the bypassed byte statistics.

If you include bypassed bytes in traffic mix views, it depicts the actual bandwidth gain achieved between the client and the server by representing the total number of optimized and unoptimized bytes exchanged on the link. Bandwidth gain statistics are lower in this view because bypassed bytes are unoptimized, using bandwidth with no corresponding caching or protocol optimization benefits.

Exclude bypassed bytes statistics in the traffic mix view, by clearing the **Include bypassed bytes** check box. This view depicts bandwidth gain on the protocols that the ProxySG intercepts and their corresponding values.

When you include or exclude bypassed bytes, only the graph data and totals are affected. The table data in the lower half of the page is not altered.

About the Default Service Statistics

The default service statistics represent bytes for traffic that has been bypassed because it did not match:

- ☐ An existing service listener
- ☐ Other rules, such as static or dynamic bypass

To view the default service bytes, click **Default Ports...** in the upper-right section of the **Statistics > Traffic Details > Traffic Mix** page.




Figure 31–3 Default Service Per Port Bytes Dialog

See "[About the Default Listener](#)" on page 112 for more information about the default service.


Viewing NetFlow Statistics

NetFlow is a network protocol developed by Cisco Systems to monitor and export IP traffic information. After you configure NetFlow on the appliance, direct the flow data to record collectors that you have already set up.

For more information on NetFlow, refer to `#(config) netflow` in the *Command Line Interface Reference*.

Viewing Traffic History

The Traffic History report shows historical data about proxies and services; you can select a particular proxy or service and then view its bandwidth usage, gain, client bytes, and server bytes over different time periods.



To view statistics for a particular proxy or service:

1. Select **Statistics > Traffic Details > Traffic History**.
2. From the **Proxy** or **Service** drop-down list, select the proxy or service of interest.
3. Select the time period you are interested in: From the **Duration** drop-down, select **Last Hour**, **Last Day**, **Last Week**, **Last Month**, or **Last Year**.
4. Click a tab (such as **BW Gain**) to display each of the four graphs for the selected proxy/service.


Graph Type	Description
<i>BW Usage</i>	Area graph showing the rate (in kilobits per second) of client, server, and bypassed traffic in the selected proxy/service during the time period
<i>BW Gain</i>	Line graph showing the bandwidth gain from optimization of the proxy/service during the time period, expressed as a multiple (for example, 2x)

<i>Client Bytes</i>	Bar graph displaying the number of bytes of the proxy / service that clients transmitted during the time period
<i>Server Bytes</i>	Bar graph displaying the number of bytes of the proxy / service that servers transmitted after optimization during the time period

- To view the average bandwidth gain and total client and server bytes for the selected proxy/service during the specified time period, look at the statistics to the left of the graph area.
- If you are interested in other time periods or proxies/services, repeat the above steps.

The graphs and statistics automatically update to reflect the time period and proxy/service you selected. Thereafter, the chart data updates automatically every 60 seconds.

Hover the mouse cursor over the chart data to view detailed values.



The colors in the report represent the following information:

- Bandwidth Usage chart:
 - Green—Client bytes
 - Blue—Server bytes
 - Brown—Bypassed bytes
 - Dark Blue—Bandwidth gain
- Bandwidth Gain chart
 - Dark Blue—Bandwidth gain
- Client and Server Byte charts:
 - Green—Intercepted client bytes
 - Blue—Intercepted server bytes
 - Brown—Bypassed bytes

Supported Proxies and Services

The **Traffic History** and **Traffic Mix** reports display data for the following proxy types and services of these proxy types.

- "Supported Proxy Types" on page 675
- "Supported Services" on page 675
- "Unsupported Proxy Types" on page 676

Supported Proxy Types

The following proxy types are supported in the **Traffic History** and **Traffic Mix** reports:

• CIFS	• Endpoint Mapper	• Flash
• FTP	• HTTP	• HTTPS Forward Proxy
• HTTPS Reverse Proxy (Only in Traffic History)	• Inbound ADN (Only in Traffic Mix)	• MAPI
• MSRPC	• QuickTime	• Real Media
• RTSP (Only in Traffic Mix)	• SSL	• TCP Tunnel
• Windows Media		

Supported Services

The following services are supported in the **Traffic History** and **Traffic Mix** reports:

• BGP	• Symantec ADN	• Symantec Management
• CIFS	• Cisco IPsec VPN	• Citrix
• Default	• Echo	• Endpoint Mapper
• FTP/FTPS	• H.323	• HTTP (External/Explicit/Internal)
• HTTPS	• IBM DS	• ICU-II
• IMAP/IMAP4S/IMAPS	• IPP	• Kerberos
• L2TP	• LDAP/LDAPS	• Lotus Notes
• LPD	• MGCP	• MMS
• MS SQL Server	• MS Terminal Services	• MySQL
• NetMeeting	• NFS	• Novell GroupWise

• Novell NCP	• Oracle/Oracle over SSL	• Other SSL
• pcAnywhere	• POP3/POP3S	• PPTP
• Print	• Remote Login Shell	• Remote Telnet
• RTMP	• RTSP (Only in Traffic Mix)	• SIP/SIP over SSL
• SMTP	• SnapMirror	• SSH
• Sybase SQL	• TACACS	• Telnet
• Time	• Tivoli DS	• VNC
• X Windows		

Note: Endpoint Mapper proxy bytes are the result of Remote Procedure Call (RPC) communication for MAPI traffic.

Unsupported Proxy Types

The **Traffic History** report does not display data for the following proxy types:


• DNS	• IM	• P2P
• SOCKS	• Telnet	
-

Viewing the Application Mix Report

The Application Mix report shows a breakdown of the Web applications running on the network. This report can give you visibility into which Web applications users are accessing, the amount of bandwidth these applications are consuming, and how much bandwidth is gained by optimization of Web applications over different time periods. The report has three parts to it:

- ❑ Line graph showing aggregated bandwidth usage or gain (see "[Viewing Bandwidth Details for Web Applications](#)" on page 678)
- ❑ Pie graph showing client/server byte distribution of Web applications (see "[Viewing Client/Server Byte Distribution for Web Applications](#)" on page 679)

- Statistical table listing client/server bytes and savings for each Web application (see "Viewing Application Statistics" on page 680)



Key:

- a: Modify the reporting time period.
- b: View client- or server byte-distribution charts and statistics.
- c: View aggregated bandwidth usage or gain graphs.
- d: Review client bytes, server bytes, and bandwidth savings.
- e: Review totals for client bytes, server bytes, and total savings.

Supported Applications

The Symantec WebFilter database contains a list of applications that it can recognize; when a user enters a URL in a Web browser, WebFilter identifies whether it is one of the supported applications. The supported applications are then included in the Application Mix report. Any URLs that are not associated with a supported application are categorized as *none*, and are included in the <Unidentified> slice in the pie chart.

Tip: To see a list of supported applications, display the **Active Sessions** report, select the **Application** filter, and look at the application names on the drop-down list. As new applications are supported, they will be updated in the WebFilter database and subsequently in the **Application** filter.

Application Reporting Requirements

Application reporting has the following requirements:

- Proxy Edition license (not a MACH5 license)
- The Symantec WebFilter feature must be enabled.
(**Configuration > Content Filtering > General**)

- ❑ A current WebFilter database must be downloaded to the ProxySG.
(**Configuration > Content Filtering > Blue Coat WebFilter**)
- ❑ The ProxySG must have one or more Web services, such as External HTTP and HTTPS, set to intercept. Bypassed Web traffic is not classified into applications.

Viewing Bandwidth Details for Web Applications

To see how much bandwidth is attributed to Web application traffic over the last hour, day, week, month, or year, view the line graph on the Application Mix report. This graph also allows you to analyze how much bandwidth you are gaining from optimization of Web applications.

To view aggregated bandwidth usage or gain statistics for Web applications:

1. Select **Statistics > Application Details > Application Mix**.
2. To see the bandwidth rate of Web applications, select the **BW Usage** tab (underneath the line graph).

The green area represents client data and the blue area represents server data.


3. To see how much bandwidth is gained due to optimization of Web applications, select the **BW Gain** tab.

The line graph indicates the bandwidth gain due to optimization of Web applications, averaged over the time interval, expressed as a multiple (for example, 2x).

4. Select the time period you are interested in from the **Duration** drop-down list.

The graphs and statistics automatically update to reflect the time period you selected. Thereafter, the chart data updates automatically every 60 seconds.

Hover the mouse cursor over the chart data to view detailed values.



The values that display when you hover the mouse cursor over the chart data, are called tool tips. These values can include:

- ❑ **C** = Client-side traffic data rate. This statistic represents the data rate calculated (to and from the client) on the client-side connection. Data rate is represented by units of bits per second (bps) from measurements that are sampled at one-minute intervals. All application protocol-level bytes are counted, including application-protocol overhead such as HTTP headers.
- ❑ **S** = Server-side traffic data rate. This statistic represents the data rate calculated (to and from the server) on the server-side connection. The data rate is represented by units of bits per second (bps) from measurements that are sampled at one-minute intervals. All application-level bytes are counted, including application overhead such as HTTP headers.

- ❑ **Unopt** = Unoptimized traffic data rate. This statistic reflects the data rate of original traffic served to/from the client or server prior to or subsequent to ADN optimization. The data rate is represented by units of bits per second (bps).
- ❑ **Opt** = Optimized traffic data rate. This statistic reflects the data rate of ADN-optimized traffic. Data rate is represented by units of bits per second (bps).
- ❑ **Gain** = Bandwidth Gain. This statistic, representing the overall bandwidth benefit achieved by object and byte caching, compression, protocol optimization, and object caching, is computed by the ratio:

`client bytes / server bytes`

and represented as a unit-less multiplication factor. Bandwidth-gain values are computed at one-minute intervals to correspond to the one-minute sampling of client and server bytes. For example, if server bytes displayed as 10kbps and client bytes was 90kbps, the bandwidth gain is represented as 9x.

- ❑ **Savings** = Bandwidth Savings. This statistic, representing the overall bandwidth savings achieved over the WAN by utilizing object and byte caching, protocol optimization, and compression, is computed by

`(client bytes - server bytes) / client bytes`

and presented as a percentage. The Savings value provides a relative percentage of bandwidth savings on the WAN link, with 100% indicating no WAN traffic at all (no server bytes) and 0% indicating that no savings were achieved by client bytes equaling server bytes. Utilizing the numbers from the above example, the equivalent savings would be $8/9 = 0.89 = 89\%$.

See Also

- ❑ ["Viewing Client/Server Byte Distribution for Web Applications"](#)
- ❑ ["Viewing Application Statistics"](#)

Viewing Client/Server Byte Distribution for Web Applications

The pie chart on the Application Mix report shows the distribution of Web applications over the last hour, day, week, month, or year. You can look at this data either by client bytes or server bytes.

To view a pie chart showing distribution of Web applications:

1. Select **Statistics > Application Details > Application Mix**.
2. Select **Client Bytes** or **Server Bytes** (tabs underneath the pie chart).
3. Select a time period from the **Duration** drop-down list.

The pie chart displays data for the seven applications with the most traffic during the selected time period. If there are more than seven applications classified during that time, the applications with the least amount of traffic are combined into an **Other** slice. The **<Unidentified>** slice includes traffic for which the URL is not a Web application, or is a Web application that is not currently supported in the database. **<Unidentified>** also includes Web traffic for applications that could not be identified because there was a problem with the WebFilter license or database.

See Also

- ❑ ["Viewing Bandwidth Details for Web Applications"](#)
- ❑ ["Viewing Application Statistics"](#)

Viewing Application Statistics

The table of statistics at the bottom of the Application Mix Report lists the following details for each Web application during the selected time period:

- ❑ Proxy Type—The name of the proxy that is handling the application.
- ❑ Client Bytes—The number of bytes (to and from the client) on the client-side connection.
- ❑ Server Bytes—The number of bytes (to and from the server) on the server-side connection.
- ❑ Savings—Bandwidth savings achieved over the WAN by utilizing object and byte caching, protocol optimization, and compression; presented as a percentage. The formula is:


$$(\text{client bytes} - \text{server bytes}) / \text{client bytes}$$

See Also

- ❑ ["Viewing Bandwidth Details for Web Applications"](#)
- ❑ ["Viewing Client/Server Byte Distribution for Web Applications"](#)

Viewing the Application History Report

The Application History report shows historical data about Web applications; you can select a particular Web application and then view its bandwidth usage, gain, client bytes, and server bytes over different time periods.



Key:

- a: View statistics for a particular Web application.**
- b: Modify the historical reporting period.**
- c: View totals for client and server bytes and the average bandwidth gain for the selected application.**
- d: Display charts for bandwidth usage, bandwidth gain, client bytes, and server bytes.**

To view statistics for a particular Web application:

1. Select **Statistics > Application Details > Application History**.
2. From the **Application** drop-down list, select the Web application of interest. This list contains any application that has been seen on the network in the last year.
3. Select the time period you are interested in: From the **Duration** drop-down, select **Last Hour**, **Last Day**, **Last Week**, **Last Month**, or **Last Year**.
4. Click a tab (such as **BW Gain**) to display each of the four graphs for the selected application.

Graph Type	Description
<i>BW Usage</i>	Area graph showing the rate (in bits per second) of client and server traffic in the selected application during the time period
<i>BW Gain</i>	Line graph showing the bandwidth gain from optimization of the application during the time period, expressed as a multiple (for example, 2x)
<i>Client Bytes</i>	Bar graph displaying the number of bytes of the application that clients transmitted during the time period

Server Bytes

Bar graph displaying the number of bytes of the application that servers transmitted during the time period

5. To view the average bandwidth gain and total client and server bytes for the selected application during the specified time period, look at the statistics to the left of the graph area.
6. If you are interested in other time periods or applications, repeat the above steps.

Viewing System Statistics

The **Statistics > System** pages enable you to view:

- ❑ "Resources Statistics" on page 682
- ❑ "Contents Statistics" on page 686
- ❑ "Event Logging Statistics" on page 689
- ❑ "Failover Statistics" on page 690

Resources Statistics


The **Resources** tabs (**CPU**, **Concurrent Users**, **Disk Use**, and **Memory Use**) allow you to view information about how the CPU, disk space and memory are being used, and how disk and memory space are allocated for cache data. You can view data allocation statistics through both the Management Console and the CLI, but disk and memory use statistics are available only through the Management Console.

Viewing CPU Utilization

Through the Management Console, you can view the average CPU utilization percentages for the ProxySG over the last hour, day, week, month, or year. You can also view CPU usage over all time periods simultaneously.

To view CPU utilization:

Select **Statistics > System > Resources > CPU**.




Note: If the ADN adaptive compression feature is enabled, the ProxySG will adjust its compression level based on its internal compression index, resulting in higher or lower CPU usage. This means that if adaptive compression is enabled, you can not rely on the CPU utilization values alone for capacity planning. Instead, you should also consider the compression index (**Statistics > ADN History > Adaptive Compression**). To determine whether adaptive compression is enabled, click the **Enable or disable adaptive compression** link.

Viewing Concurrent Users

The **Concurrent Users** tab shows users (IP addresses) that are being intercepted by the ProxySG. The duration intervals that you can view concurrent use are for the last hour, day, week, month, and year. Only unique IP addresses of connections intercepted by proxy services are counted toward the user limit.

To view concurrent users:

Click **Statistics > System > Resources > Concurrent Users**.



Viewing Disk Use Statistics

The **Disk Use** tab shows statistics about the ProxySG appliance disk usage.


- ❑ **System objects**—Percentage of storage resources currently used for system objects.
- ❑ **Access logs**—Percentage of storage resources currently used for access logs.
- ❑ **Cache**—Percentage of storage resources available for cache objects. This statistic represents both cache that is in use and the remaining space for cache.

The total disk usage is the sum of the first two statistics: system objects usage and the access logs usage. SNMP monitoring reports on this total for disk usage alerts; for more information on SNMP monitoring, see [Section D: "Monitoring Network Devices \(SNMP\)"](#) on page 1319.

The total disk installed is the sum of all three statistics: system objects usage, access logs usage, and available cache.

To view disk use statistics:

Select **Statistics > System > Resources > Disk Use**.



Viewing Memory Use Statistics

The **Memory Use** tab shows the absolute values and percentages of RAM being used. The fields on the Memory Use tab are:


- ❑ **Committed by system**—RAM required for operating system.
- ❑ **Committed by applications**—RAM required by system applications.
- ❑ **Committed cache buffers**—RAM has been allocated and is still in use.
- ❑ **Reclaimable cache buffers**—Set of memory segments used to cache object data and accelerate performance. RAM has retention value; cache buffers contain useful data.

Note: The Kernel attempts to maximize the number and lifetime of cache buffers, but if needed, it will recover cache buffers using the LRU replacement algorithm to satisfy a memory allocation request.

- ❑ **Free**—RAM that has no retention value.

To view memory use statistics:

Select **Statistics > System > Resources > Memory Use**.



Contents Statistics

The **Contents** tabs (**Distribution** and **Data**) allow you to see information about objects currently stored or served that are organized by size. The cache contents include all objects currently stored by the ProxySG appliance. The cache contents are not cleared when the appliance is powered off.


Viewing Cached Objects by Size

The **Distribution** tab shows the cached objects currently stored by the ProxySG appliance and their size.

- ❑ **Use Logarithmic Scale**— Enables all cached objects with a wide range of values to be represented in the graph. For example, the ProxySG appliance might have one million cached objects of 1KB or less in size and only 10 objects of 500kb or less in size. If the logarithmic scale is disabled, larger objects might not be visible on the graph.

To view the distribution of cache contents:

Select **Statistics > System > Contents > Distribution**.




Viewing the Number of Objects Served by Size

The **Data** tab displays the number of objects served by the ProxySG appliance that are organized by size. The chart shows you how many objects of various sizes have been served.

☐ **Objects in Cache**—The number of objects that are currently cached

To view the number of objects served:

Select **Statistics > System > Contents > Data**.




Event Logging Statistics

The event log contains all events that have occurred on the ProxySG. Configure the level of detail available by selecting **Maintenance > Event Logging > Level** (For details, see ["Selecting Which Events to View"](#) on page 1309 in Chapter 2).

To view the event log:

1. Select **Statistics > System > Event Logging**.



2. Click **Log start** or **Log end** or the forward and back arrow buttons to move through the event list.
3. (Optional) Click the **Poll for new events** check box to poll for new events that occurred while the log was being displayed.


Note: The Event Log cannot be cleared.

Failover Statistics

At any time, you can view statistics for any failover group you have configured on your system.

To view failover statistics:

1. Select **Statistics > System > Failover**.



Status	
Failover Group:	10.9.16.150
Failover status:	
Multicast address:	224.1.2.3
Local address:	10.9.16.150
State:	MASTER
Flags:	R (Real IP)

2. From the **Failover Group** drop-down list, select the group to view.

The information displayed includes the multicast address, the local address, the state, and any flags, where **V** indicates that the group name is a virtual IP address, **R** indicates that the group name is a physical IP address, and **M** indicates that this machine can be configured to be the master if it is available.

Active Sessions—Viewing Per-Connection Statistics

Viewing active sessions enables you to view detailed statistics about proxied sessions, errored sessions, bypassed connections, and ADN inbound connections.

- Viewing the proxied sessions provides information for diagnostic purposes.
- Viewing bypassed connections helps identify new types of traffic flowing through the ProxySG appliance, as well as traffic flows that would benefit from optimization.
- Viewing active ADN inbound connections provides information for diagnostic purposes.
- Viewing errored sessions enables you to track details for troubleshooting.

For specific information, see ["Analyzing Proxied Sessions"](#) on page 691, ["Analyzing Bypassed Connections Statistics"](#) on page 703, and ["Viewing Errored Sessions and Connections"](#) on page 705.

Note: You can also view session statistics for ADN inbound connections, which is described in ["Reviewing ADN Active Sessions"](#) on page 759.

See Also

- ❑ ["Example Scenarios Using Active Sessions for Troubleshooting"](#) on page 691
- ❑ ["Analyzing Proxied Sessions"](#) on page 691

- ❑ "Analyzing Bypassed Connections Statistics" on page 703
- ❑ "Viewing Errored Sessions and Connections" on page 705

Example Scenarios Using Active Sessions for Troubleshooting

An administrator is setting up a Common Internet File System (CIFS) over ADN and the CIFS does not appear to be working. The administrator can use the Active Sessions feature on the ProxySG to filter for any CIFS sessions that produced an error. If the ProxySG did not report an error, the administrator still has some information about the session that can help diagnose the failure without the use of a packet capture.

The following list describes two other examples when using active sessions can help with troubleshooting problems.

- ❑ A site-wide problem is occurring and the administrator uses active sessions to diagnose the failure. If it is a problem with DNS, for example, there will be a large number of sessions with DNS errors.
- ❑ In protocols where errors might not be communicated another way (such as CIFS, TCP, or tunnels), active sessions record the actual error.

Analyzing Proxied Sessions

The **Statistics > Active Sessions > Proxied Sessions** page provides an immediate picture of the sessions and the protocol types, services, bytes, savings, and other statistics. These statistics are derived from WAN optimization and object caching and are associated with client traffic.

The first time you view the **Proxied Sessions** page, no data is displayed. To display proxied sessions data, click **Show**. The statistics displayed in the window are not automatically updated. To update the statistics, click **Show** again.

The screenshot shows the 'Proxied Sessions' tab selected. The 'Filter' is set to 'None' and the 'Show' button is highlighted. Below the filter, there are checkboxes for 'Display the most recent' and 'Show errored sessions only'. The table below shows the following data:

Client	Server	A	S	FW	I	Duration	Client Bytes	Server Bytes	
10.9.44.212:3612	m1.2mdn.net:80					4.6 min	31,388	0	
10.9.44.212:3632	ad.trafficmp.com:80					1.9 min	9,511	9,553	
10.9.44.212:3632	ad.trafficmp.com:80					1.0 min	1,280	1,276	

Below the table are buttons for 'Terminate Session' and 'Download'. At the bottom, it says 'Total displayed sessions: 25 Total displayed connections: 52'.

Important: Use the statistics on the **Proxied Sessions** pages as a diagnostic tool only. The **Proxied Sessions** pages do not display every connection running through the ProxySG. This feature displays only the *active* sessions—one client connection (or several), together with the relevant information collected from other connections associated with that client connection. Because it displays only *open* connections, you cannot use the data for reporting purposes.

The **Proxied Sessions** page displays statistics for the following proxies:

- ☐ Adobe HDS
- ☐ Adobe HLS
- ☐ CIFS
- ☐ Endpoint Mapper
- ☐ Flash
- ☐ FTP
- ☐ HTTP
- ☐ HTTPS Forward Proxy
- ☐ HTTPS Reverse Proxy
- ☐ MAPI
- ☐ MSRPC
- ☐ MS Smooth
- ☐ QuickTime
- ☐ Real Media
- ☐ SSL
- ☐ STunnel
- ☐ TCP Tunnel
- ☐ Websocket (Introduced in SGOS 6.5.5.7)
- ☐ Windows Media

Viewing Proxied Sessions

Client connections are available for viewing as soon as the connection request is received. However, if delayed intercept is enabled, the connection is not shown until the three-way handshake completes. Server connections are registered and shown in the table after the `connect` call completes.

To view proxied sessions:

1. Select the **Statistics > Sessions > Active Sessions > Proxied Sessions** tab.
2. Select a filter from the **Filter** drop-down list.

Important: It is important to select a filter before clicking **Show** to minimize the time it might take for a busy ProxySG to download the list of active sessions.

3. Enter the appropriate information for the filter you have selected:

Filter	Information to Enter
<i>Application</i> (For Proxy Edition license only)	Select a Web application from the drop-down list. All supported applications appear on this list; this list will automatically populate with new applications as they are added to the WebFilter database. (Note that this requires that your system downloads an updated WebFilter database; by default, your system will automatically check for updates.)
<i>Client Address</i>	Enter the client's IP address or IP address and subnet mask
<i>Client Port</i>	Enter a client port number.
<i>ICAP</i> (For Proxy Edition license only)	Select the ICAP service type from the drop-down list: Any, REQMOD, RESPMOD Select the service name from the Service drop-down list. Select the ICAP state from the Status drop-down list: Any, transferring, deferred, scanning, completed Notes: - The ICAP filtering fields are optional. If you leave all the options set to Any, all ICAP-enabled sessions are listed. - To see entries that represents a session instead of a connection, you must expand that row (by clicking the Client column) to see all the connections inside the session.
<i>Proxy</i>	Select a filter from the drop-down list.
<i>Server Address</i>	Enter the IP address or hostname of the server. Hostname filters automatically search for suffix matches. For example, if you filter for example.com, test.example.com is included in the results.
<i>Server Port</i>	Enter a server port number.
<i>Service</i>	Select an enabled service from the drop-down list.

4. (Optional) To limit the number of connections to view, select **Display the most recent** and enter a number in the results field. This optimizes performance when there is a large number of connections.
5. (Optional) To view the current errored proxied sessions, select **Show errored sessions only**. For more details, see "[Viewing Errored Sessions and Connections](#)" on page 705.

6. Click **Show**.

Downloading Proxied Session Statistics

To save and share session statistics data for diagnostic purposes, you can download the current proxied sessions statistics and save them in an Excel file.

To download proxied session statistics:

1. Click **Download**. The Save dialog displays.
2. Navigate to the location to save the text file and click **Save**. The text file contains all the statistics for the current proxied sessions.
3. (Optional) Save the data in an Excel file by copying the contents of the text file, opening Excel, and selecting **Edit > Paste Special**.

Terminating a Proxied Session

Terminating an active session causes any operation in progress on the session to be interrupted, so it is not advised to do so unless there is a specific condition that needs to be remedied. When you terminate a proxied session, the ProxySG terminates both the client-side and server-side connections.

For example, a CIFS session might report an error that was preventing it from being accelerated. The administrator would then reconfigure some settings on the client or server to fix the problem. After that, the administrator could terminate the session on the ProxySG, which would force the client to connect again and allow the new connection to be accelerated.

To terminate a proxied session:

Select the session in the list and click **Terminate Session**.

About the Proxied Sessions Statistics

When reviewing the proxied session statistics, note that:

- ❑ Active client and server connections are displayed in black.
- ❑ Inactive connections are displayed in gray.
- ❑ Errored connections are displayed in red (when you select the **Show errored sessions only** check box).
- ❑ Session and connection totals are displayed on the bottom left side of the page.

The following table describes the per-column statistics and the various icons on the **Proxied Sessions** page.

Table 31–1 Column and Icon Descriptions on the Proxied Sessions Page

Column or Icon	Description
<i>Client</i>	IP address and port of the client PC (or other downstream host). When the client connection is inactive, the contents of this column are unavailable (gray). A client connection can become inactive if, for example, a client requests a large object and then aborts the download before the ProxySG has finished downloading it into its cache. When the session has multiple client connections, a tree view is provided. See "Viewing Sessions with Multiple Connections" on page 700 for more information.
<i>Server</i>	Final destination of the request. By default, the hostname is displayed. However, if a user entered an IP address in the URL, the IP address is displayed. The contents of this column are unavailable if the server connection is inactive. This can occur when a download has completed (and the server connection is closed or returned to the idle pool), but the object is still being served to the client. If a server connection was never made (a pure cache hit case), the Server column displays the hostname (or IP address) of the requested server. Active server connections are shown in black; inactive connections are gray.
A  	ADN. Indicates that the server connection is flowing over an ADN tunnel. If the icon does not display, it indicates that an ADN tunnel is not in use. Encrypted ADN tunnel.
S 	SOCKS. Indicates that the upstream connection is being sent through a SOCKS gateway. If the icon does not display, it indicates that a SOCKS gateway is not in use.
FW 	Forwarding. Indicates that the upstream connection is being sent through a forwarding host. If the icon does not display, it indicates that forwarding is not in use.

Table 31–1 Column and Icon Descriptions on the Proxied Sessions Page (Continued)

Column or Icon	Description
/      	ICAP services are displayed if you have a Proxy Edition license only. Indicates an ICAP-enabled session. If the icon does not display, ICAP is not supported for that session. Different icons are used to indicate the ICAP state of the session. - Transferring (arrow) — ICAP requests are being transferred to the ICAP server - Deferred (clock) — ICAP scanning requests have been deferred until the full object has been received - Scanning (magnifying glass) — ICAP requests are in the process of being scanned - Completed (checkmark) — ICAP scanning requests completed successfully - Inactive (i) — The ICAP feature is inactive for the session - Unsupported (no icon) — ICAP is not supported for the corresponding session
<i>Duration</i>	Displays the amount of time the session has been established.
<i>Client Bytes</i>	Represents the number of bytes (to and from the client) at the socket level on the client connection. All application-level bytes are counted, including application overhead such as HTTP headers, CIFS headers, and so on. TCP and IP headers, packet retransmissions, and duplicate packets are not counted. See "About the Byte Totals" on page 701 for more information.
<i>Server Bytes</i>	Represents the number of bytes (to and from the server) at the socket level on the server connection. All application-level bytes are counted, including application overhead such as HTTP headers, CIFS headers, and so on. If the traffic is flowing through an ADN tunnel, the bytes are counted after ADN optimization, meaning that compressed byte counts are displayed. TCP and IP headers, packet retransmissions, and duplicate packets are not counted. See "About the Byte Totals" on page 701 for more information.
<i>Savings</i>	Displays the bandwidth gain for the session and the savings in bandwidth. When the request results in a pure cache hit, this column displays 100%.

Table 31–1 Column and Icon Descriptions on the Proxied Sessions Page (Continued)

Column or Icon	Description
<i>C</i> 	Compression. When displayed in color, this icon indicates that an ADN Tunnel is in use and <code>gzip</code> compression is active in either direction on that tunnel. This icon has three states: • Active (color icon) • Inactive (gray icon) • Not possible (not displayed)
<i>BC</i> 	Byte Caching. When displayed in color, this icon indicates that an ADN Tunnel is in use and byte caching is active in either direction on that tunnel. This icon has three states: • Active (color icon) • Inactive (gray icon) • Not possible (not displayed) Note: If the control connection fails to establish, the two ADN peers cannot synchronize their byte cache dictionaries. This can happen, for example, in a transparent unmanaged ADN if the concentrator peer sends a control IP address that is not accessible from the branch peer. When a control connection with the peer is not established and the dictionaries are out of sync, a warning icon  displays in the BC column to alert you of the problem. Although byte caching is enabled, it is not in use. If you see this icon, you can fix the issue by specifying preferred IP addresses on the concentrator. See the preferred-ip-addresses command in the <i>SGOS Command Line Reference Guide</i> for more information.
<i>OC</i> 	Object Caching. When displayed in color, this icon indicates that an HTTP, HTTPS, CIFS, Streaming, or FTP proxy is in use and the content is cacheable. This icon has three states: • Active (color icon) • Inactive (gray icon) • Not possible (not displayed) The icon: • Is unavailable if the content is non-cacheable (or for CIFS, when the entire connection is non-cacheable—not on an object-by-object basis). • Is not displayed for MAPI and TCP-Tunnel traffic. • Does not indicate a cache hit; it indicates only that the object is cacheable.

Table 31–1 Column and Icon Descriptions on the Proxied Sessions Page (Continued)

Column or Icon	Description
<i>P</i> 	Protocol Optimization. When displayed in color, this icon indicates that a proxy is in use that is capable of performing latency optimizations. These proxies include HTTP, HTTPS, CIFS, MAPI, MMS and RTSP. This icon has three states: • Active (color icon) • Inactive (gray icon) • Not possible (not displayed)
<i>BM</i> 	Bandwidth Management. When displayed in color, this icon indicates that either the client or server connection has been assigned to a bandwidth class. This icon has two states: • Active (color icon) • Inactive (gray icon)
<i>E</i> 	Encryption. When displayed in color, this icon indicates that an ADN Tunnel is in use and encryption is active in either direction on that tunnel. This icon has three states: • Active (color icon) • Inactive (gray icon) • Not possible (not displayed)
<i>Service Name</i>	Displays the service used by the session. Even if a client connection is handed off to a different application proxy, this column shows the service name of the original service that intercepted the client connection.
<i>Application (For Proxy Edition license only)</i>	Displays the name of the Web application used by the session. If no application is listed, the session is either not a Web application, the application is not currently supported, or the appliance does not have a valid WebFilter license.
<i>Protocol</i>	Displays the protocol used by the session.

Table 31–1 Column and Icon Descriptions on the Proxied Sessions Page (Continued)

Column or Icon	Description
<i>Detail</i>	Provides additional information. For example, it can indicate that a CIFS connection is "pass-through" due to SMB signing or "Thin client processing enabled" for a connection. The Detail column also displays the following errors: • Errors connecting upstream (TCP errors, ADN network errors) • Unexpected network errors after connecting (e.g., read errors) • Request-handling errors (parse errors, unknown method or protocol, unsupported feature) • Response-handling errors (parse errors, unknown method or protocol, unsupported feature, unexpected responses such as HTTP 500 errors from OCS) • Unexpected internal errors • DNS errors and DNS resolve failures • External service errors such as ICAP, BCAA, and so on See "Viewing Errored Sessions and Connections" on page 705.

Viewing Additional Information

Place the cursor over the following components or fields to get more information:

- ❑ Table column headers—Displays the full name of the column header.
- ❑ Row values.
- ❑ Acceleration icons (**C**, **BC**, **OC**, **P**, **BM**)—Displays the icon identity.
- ❑ ADN, SOCKS, and FW icons—Displays the upstream host of that type being communicated with, if any.
- ❑ ICAP icons—Displays the type of service (REQMOD and/or RESPMOD), the name of the service, and the session's ICAP state (transferring, deferred, scanning, or completed).
- ❑ Client—Displays the full hostname or IP address.
- ❑ Server—Displays the client-supplied destination IP address, the destination server address (the final server address to which the proxy is connecting), and when available, the address of the upstream forwarding host and the address of the upstream SOCKS gateway.

About MMS Streaming Connections

The Active Sessions feature displays connection statistics for MMS streams over HTTP, TCP, or UDP only. Multicast connections are not displayed. When an MMS stream is displayed, the service name is listed as **HTTP** or **MMS** (depending on the transport used) and the protocol indicates Windows Media.

10.9.59.48:2597	msent.wmod.llnwd.net:80	14 sec	494,885	166,012

Figure 31–4 MMS Streaming Connection Example

Viewing Sessions with Multiple Connections

When multiple client or server connections are associated with a single session, the **Client** column provides a tree-view that allows you to expand the row to view more details about the associated connections. The tree view is represented by the  icon.

The following figure shows an HTTP example of this tree view.

10.9.59.48.3579	amch.questionmarket.co...	1 sec	13,760	4,496	206%			HTTP	HTTP
10.9.59.48.3579	amch.questionmarket.co...	1 sec	13,760	1,638	740%			HTTP	HTTP
0	i.cnn.net:0	0 sec	0	0				HTTP	HTTP
0	ads.cnn.com:80	0 sec	0	1,508				HTTP	HTTP
0	view.atdmt.com:80	0 sec	0	700				HTTP	HTTP
0	ad.doubleclick.net:80	0 sec	0	650				HTTP	HTTP

Figure 31–5 Multiple Server Connections Example

HTTP

The tree view displays (as shown above) for HTTP if multiple hosts are contacted during a session or if pipelining is used.

FTP

FTP uses multiple, concurrent connections. These are represented as separate rows in the tree view, as shown in the following figure.

Client	Server	A	S	FW	I	Duration	Client Bytes	Server Bytes	Savings	C	BC	OC	P	BM	Service Name
10.9.44.163:4939	130.14.29.30:21					1.6 min	1,401,203	702,418	49.35%						FTP
10.9.44.163:4939	130.14.29.30:21					1.6 min	2,108	2,203	0%						FTP
10.9.44.163:4959						6 sec	698,880	0	100%						FTP
10.9.44.163:4939	130.14.29.30:50149					17 sec	698,880	698,880	0%						FTP
10.9.44.163:4939	130.14.29.30:50303					1 sec	879	879	0%						FTP
10.9.44.163:4939	130.14.29.30:50122					0 sec	456	456	0%						FTP

Figure 31–6 FTP Connections Example

CIFS, MAPI, and Endpoint Mapper do not display multiple connections.

MMS

The active sessions feature displays MMS streams that have a client associated with them. MMS streams that do not have a client associated with them (multicast, content management requests, and so on) are not displayed. MMS streams are displayed as follows:

- ❑ MMS UDP streams have two connections, one for data and one for control.
- ❑ MMS TCP streams have a single connection.
- ❑ MMS HTTP streams have a single connection.

For additional information about streaming connections, see "[About MMS Streaming Connections](#)" on page 699.

Expanding the Active Sessions Tree View

When expanded, the tree view displays per-connection statistics for the session, as shown in the following example. To expand the results for a connection, click the arrow to the left of the client IP address.

Client	Server	A	S	FW	I	Duration	Client Bytes	Server Bytes	Savings
10.9.44.163:1902	m1.2mdn.net:80					3.7 min	257,528	155,555	39.76%
10.9.44.163:1902	m1.2mdn.net:80					3.7 min	8,606	9,433	0%
:0	m1.2mdn.net:80					0 sec	51,967	0	100%

Figure 31–7 Active Sessions Tree View (Expanded)

The **Savings** column result differs according to the server or client byte totals:

- ☐ Zero client bytes: displays no savings.
- ☐ Zero client and server bytes: displays no savings.
- ☐ Client and server are greater than zero: displays the calculated savings.

About the Byte Totals

The client and server byte total is the sum of all bytes going to and from the client or server. All application-level bytes are counted, including application overhead such as HTTP headers, CIFS headers, and so on. TCP and IP headers, packet retransmissions, and duplicate packets are not counted.

The following sections describe some of the factors that can affect the byte totals.

ADN Tunnels

If the traffic is flowing through an ADN tunnel, the bytes are counted after ADN optimization, meaning that compressed byte counts are displayed.

Multiple Server Connections

A single client connection can use many server connections. The server byte counts include the total bytes transferred over all server connections accessed over the lifetime of a client connection. Even though a server connection can serve many clients, the same server byte is never included in more than one client connection total.

Aborted Downloads

In some cases, you might see the server bytes increasing even after the client has closed the connection. This can occur when a client requests a large object and aborts the download before receiving the entire object. The server bytes continue to increase because the ProxySG is retrieving the object for caching. You can change this behavior by enabling the bandwidth gain mode.

To enable the bandwidth gain mode:

1. Select **Configuration > Proxy Settings > HTTP Proxy > Acceleration Profile**.
2. Select **Enable bandwidth gain mode**.
3. Click **Apply**.

An alternative way to do this is to add the following to policy:

```
<cache>
delete_on_abandonment (yes)
```

Explicit Proxying and Pipelining

If clients are explicitly proxied and the session has multiple connections or is pipelined, no client bytes are displayed and the expanded server connections display no savings when the tree view is shown. This is because the ProxySG is downloading the content before serving it to the client.

What Is Not Displayed

The **Proxied Sessions** page does not display statistics for:

- ❑ IM (Yahoo, AOL, MSN), DNS, SOCKS, and Telnet
- ❑ Inbound ADN connections (These display on the **ADN Inbound Connections** page.)
- ❑ Bridged connections
- ❑ Administrative connections (Management Console, SSH console, SNMP, DSAT, access-logging, Director, and so on)
- ❑ Off-box processing connections (ICAP, WebPulse, and so on)

Note: In some cases, an administrative or off-box connection might correspond to a specific client connection, for example, an ICAP AV scanning connection associated with a specific HTTP client connection. However, the byte counts collected from the ICAP AV scanning connection are not included in the Active Sessions display.

Viewing HTML and XML Views of Proxied Sessions Data

Access the following URLs to get HTML and XML views of active session statistics:

- ❑ **HTML:** https://ProxySG_IP:8082/AS/Sessions/
- ❑ **XML:** https://ProxySG_IP:8082/AS/ProxiedConnections/xml

See Also

- ❑ ["Analyzing Bypassed Connections Statistics"](#)
- ❑ ["Viewing Errored Sessions and Connections"](#)

Analyzing Bypassed Connections Statistics

The **Statistics > Sessions > Active Sessions > Bypassed Connections** page displays data for all unintercepted TCP traffic.

When the appliance is first installed in an in-path deployment, all services are bypassed by default. By analyzing the connection data in the **Bypassed Connections** page, you can review the types of traffic flowing through the appliance to identify traffic flows that would benefit from optimization. The **Bypassed Connections** page is also useful for identifying new types of traffic flowing through the appliance.

Viewing, Downloading, and Terminating Bypassed Connections

The **Bypassed Connections** page displays data for connections that were not intercepted due to one of the following:

- ❑ A service has not been configured to intercept the traffic.
- ❑ A static or dynamic bypass rule caused the traffic to be bypassed.
- ❑ The interface transparent interception setting is disabled.
- ❑ Restrict intercept is configured.

To view bypassed connections:

1. Select **Statistics > Sessions > Active Sessions > Bypassed Connections**.
2. Select a filter from the **Filter** drop-down list.

Important: It is important to select a filter before clicking **Show** to minimize the time it might take for a busy ProxySG to download the list of active sessions.

3. Enter the appropriate information for the filter you have selected:

Filter	Information to Enter
<i>Client Address</i>	Enter the client's IP address or IP address and subnet mask
<i>Client Port</i>	Enter a client port number.
<i>Server Address</i>	Enter the IP address or hostname of the server. Hostname filters automatically search for suffix matches. For example, if you filter for example.com, test.example.com is included in the results.
<i>Server Port</i>	Enter a server port number.
<i>Service</i>	Select an enabled service from the drop-down list.

4. (Optional) To limit the number of connects to view, select **Display the most recent** and enter a number in the results field. This helps optimize performance when there is a large number of connections.

5. (Optional) To view the current errored bypassed connections, select **Show errored sessions only**. For more details, see ["Viewing Errored Sessions and Connections"](#) on page 705.
6. Click **Show**.

Note the following:

- ❑ Unavailable connections (gray) indicate connections that are now closed.
- ❑ Previously-established connections displayed with (<--?-->) text indicate that the direction of these connections is unknown.
- ❑ One-way connections are displayed in color.

To download bypassed connections statistics:

1. Click **Download**. The Save dialog displays.
2. Navigate to the location to save the text file and click **Save**. The text file contains all the statistics for the current bypassed connections.
3. (Optional) Save the data in an Excel file by copying the contents of the text file, opening Excel, and selecting **Edit > Paste Special**.

To terminate a bypassed connection:

Select a connection in the list and click **Terminate Connection**.

About Bypassed Connection Statistics

The following table describes the column headings on the **Bypassed Connections** page.

Table 31–2 Table Column Heading Descriptions on the Bypassed Connections Page

Column Heading	Description
<i>Client</i>	IP address and port of the client PC (or other downstream host).
<i>Server</i>	Server IP address and port number.
<i>Duration</i>	Displays the amount of time the connection has been established.
<i>Bypassed Bytes</i>	Displays the total number of bypassed bytes for the connection.
<i>Service Name</i>	Displays the service used by the connection.
<i>Details</i>	Provides additional information. For example: • One-way traffic (forward) • One-way traffic (reverse) • Previously established • Bypassed because of network interface setting

Note: SGOS 5.3 and later bypasses CIFS sessions that require message signing or server signatures. Object caching and protocol optimization are inactive for these CIFS sessions, and the message in the Details field is **Server requires security signatures**.

Viewing HTML and XML Views of Bypassed Connections Data

Access the following URLs to get HTML and XML views of active session statistics:

- ❑ **HTML:** `https://ProxySG_IP:8082/AS/BypassedConnections/`
- ❑ **XML:** `https://ProxySG_IP:8082/AS/BypassedConnections/xml`

See Also

- ❑ "Active Sessions—Viewing Per-Connection Statistics"
- ❑ "Example Scenarios Using Active Sessions for Troubleshooting"
- ❑ "About the Proxied Sessions Statistics"
- ❑ "Analyzing Proxied Sessions"
- ❑ "Viewing Errored Sessions and Connections"

Viewing Errored Sessions and Connections

Although you can view current errored sessions on the **Proxied Sessions**, **Bypassed Connections**, and **ADN Inbound Connections** pages by selecting a check box, you can also view both current and historical errored sessions on the **Statistics > Sessions > Errored Sessions** pages. There are three pages: one for errored proxied sessions, one for errored bypassed connections, and one for ADN inbound connections.

The **Detail** column displays the type of error received. For example, if you open a browser and enter a URL for which the hostname cannot be resolved, the information displayed in the Detail column is **DNS error: unresolved hostname (Network Error)**.

To view errored sessions or connections:

1. Select **Statistics > Sessions > Errored Sessions**. Select the **Proxied Sessions** page, the **Bypassed Connections** page, or the **ADN Inbound Connections** page, depending on the type of Errored sessions you want to view.
2. Select a filter from the **Filter** drop-down list.
3. Enter the appropriate information for the filter you have selected:

Filter

*Application
(For Proxy Edition
license only)*

Client Address


Information to Enter

Select the Web application from the drop-down list. All supported applications appear on this list.

Enter the IP address of client.

<i>Client Port</i>	Enter a client port number. Client port is not available for ADN inbound connections.
<i>ICAP</i> (For Proxy Edition license only)	Select the type of service from the drop-down list: Any, REQMOD, RESPMOD Select the service name from the Service drop-down list. Select the ICAP state from the Status drop-down list: Any, transferring, deferred, scanning, completed Note: The ICAP filtering fields are optional. If you leave all the options set to Any , all ICAP-enabled sessions will be listed. The ICAP filter is available for proxied sessions only.
<i>Proxy</i>	Select a proxy from the drop-down list. Proxy filter is available for proxied sessions only.
<i>Server Address</i>	Enter the IP address of server.
<i>Server Port</i>	Enter a server port number.
<i>Service</i>	Select a service from the drop-down list. Service is not available for ADN inbound connections.
<i>Peer Address</i>	Enter the IP address of peer. Peer address is available for ADN inbound connections only.

4. (Optional) To limit the number of connections to view, select **Display the most recent** and enter a number in the results field.
5. Click **Show**.
6. Scroll to the right to display the **Detail** column and view error details. To sort by error type, click the **Detail** column header. The **Age** column displays how long it has been since that session ended.



BC	OC	P	BM	Service Name	Protocol	Detail	Age
				HTTP	HTTP	DNS error: server failure (Network Error)	0 sec
				HTTP	HTTP	DNS error: server failure (Network Error)	0 sec
				HTTP	HTTP	DNS error: server failure (Network Error)	0 sec
				HTTP	HTTP	Invalid request (Request Error)	0 sec
				HTTP	HTTP	TCP error: operation failed (Network Error)	1.46 Days

Figure 31–8 Errored Connections Details

See "About the Proxied Sessions Statistics" on page 694 for descriptions of each column and icon in the Errored Sessions pages.

To terminate an errored session or connection:

Select an errored session or connection in the list and click **Terminate Session** (for proxied Errored sessions) or **Terminate Connection** (for bypassed errored connections and ADN inbound connections).

Downloading Errored Sessions or Connections Statistics

For troubleshooting purposes, you can download errored session (proxied) or errored connection (bypassed or ADN-inbound) statistics and save the data in an Excel file.

To download errored sessions or connections statistics:

1. Click **Download**. The Save dialog displays.
2. Navigate to the location to save the text file and click **Save**. The text file contains all the statistics for the errored sessions.
3. (Optional) Save the data in an Excel file by copying the contents of the text file, opening Excel, and selecting **Edit > Paste Special**.

See Also

- ❑ ["Active Sessions—Viewing Per-Connection Statistics"](#)
- ❑ ["Example Scenarios Using Active Sessions for Troubleshooting"](#)
- ❑ ["Analyzing Proxied Sessions"](#)
- ❑ ["About the Proxied Sessions Statistics"](#)
- ❑ ["Analyzing Bypassed Connections Statistics"](#)
- ❑ ["Reviewing ADN Active Sessions"](#)

ADN History

The **Statistics > ADN History** pages allow you to view either usage statistics or gain statistics and either unoptimized bytes or optimized bytes through the ADN History tab. For more information about these statistics, see ["Reviewing ADN History"](#) on page 758.

Bandwidth Management Statistics

The **Statistics > Bandwidth Mgmt** pages display the current class and total class statistics. See ["Bandwidth Management Statistics"](#) on page 604 for more information about these statistics.

SG Client Statistics

The **Statistics > SG Client History** pages display the SG Client Manager statistics. Refer to the *ProxyClient Configuration and Deployment Guide* for more information about these statistics.

Network Interface History Statistics

The **Statistics > Network > Interface History** page displays the traffic to and from each interface, including VLAN traffic, on the ProxySG. See ["Viewing Interface Statistics"](#) on page 1248 for more information.

WCCP Statistics

The **Statistics > Network > WCCP** page displays whether WCCP is enabled and displays the number of packets redirected by the ProxySG, status of the configured service groups including details on the **Here I am, I see you** and the number of **redirect assign** messages sent to the routers in the group by the ProxySG. See ["Viewing WCCP Statistics and Service Group Status"](#) on page 786 for more information.

Protocol Statistics

The **Statistics > Protocol Details** pages provide statistics for the protocols serviced by the ProxySG. These statistics should be used to compliment the statistics in the **Traffic History** and **Traffic Mix** pages.

The descriptions of these statistics are located in the proxy services to which they pertain. The following list provides a listing of these statistics and describes where to find additional information.

- ❑ CIFS History

The **Statistics > Protocol Details > CIFS History** pages enable you view statistics for CIFS objects, CIFS bytes read, CIFS bytes written, and CIFS clients. See ["Reviewing CIFS Protocol Statistics"](#) on page 268 for more information about these statistics.

- ❑ HTTP/FTP History

The **Statistics > Protocol Details > HTTP/FTP History** pages enable you view statistics for HTTP/HTTPS/FTP objects, HTTP/HTTPS/FTP bytes, HTTP/HTTPS/FTP clients, client compression gain, and server compression gain. See ["Viewing FTP Statistics"](#) on page 299 and ["Understanding HTTP Compression"](#) on page 193 for more information about these statistics.

For HTTP/FTP bandwidth usage statistics, see the **Traffic Mix** and **Traffic History** pages.

- ❑ IM History

The **Statistics > Protocol Details > IM History** pages enable you view statistics for IM connection data, IM activity data, and IM clients.

- ❑ MAPI History

The **Statistics > Protocol Details > MAPI History** pages enable you view statistics for MAPI client bytes read, MAPI client bytes written, and MAPI clients. See ["Reviewing Endpoint Mapper Proxy Statistics"](#) on page 279 for more information about these statistics.

For MAPI bandwidth usage statistics, see the **Traffic Mix** and **Traffic History** pages.

- ❑ P2P History

The **Statistics > Protocol Details > P2P History** pages enable you view statistics for P2P data, P2P clients, and P2P bytes. Refer to the P2P information in the *Visual Policy Manager Reference* for more information about these statistics.

- ❑ Shell History

The **Statistics > Protocol Details > Shell History** pages enable you view statistics for shell clients. See "[Viewing Shell History Statistics](#)" on page 317 for more information about these statistics.

- ❑ SOCKS History

The **Statistics > Protocol Details > SOCKS History** pages enable you view statistics for SOCKS clients, SOCKS connections, client compression gain, and server compression gain. See "[Viewing SOCKS History Statistics](#)" on page 308 for more information about these statistics.

- ❑ SSL History

The **Statistics > Protocol Details > SSL History** pages enable you view statistics for unintercepted SSL data, unintercepted SSL clients, and unintercepted SSL bytes. See "[Viewing SSL History Statistics](#)" on page 232 for more information about these statistics.

- ❑ Streaming History

The **Statistics > Protocol Details > Streaming History** pages enable you view statistics for Windows Media, Real Media, QuickTime, current streaming data, total streaming data, and bandwidth gain. See "[Viewing Streaming History Statistics](#)" on page 562 for more information about these statistics.

For MMS bandwidth usage statistics, see the **Traffic Mix** and **Traffic History** pages.

Health Monitoring Statistics

The **Statistics > Health Monitoring** page enables you to get more details about the current state of the health monitoring metrics. Health monitoring tracks the aggregate health of the ProxySG and aids in focusing attention, if the health state changes. See [Chapter 72: "Monitoring the ProxySG"](#) on page 1299 for information about health monitoring.

Health Check Statistics

Use the **Statistics > Health Checks** page to view the state of various health checks: whether the health check is enabled or disabled, if it is reporting the device or service to be healthy or sick, or if errors are being reported. See [Chapter 73: "Verifying the Health of Services Configured on the ProxySG"](#) on page 1353 for more information.

Access Logging

The **Statistics > Access Logging** pages enable you to view the log tail, log size, and upload status of the access log. See "[Viewing Access-Log Statistics](#)" on page 623 for more information.

Advanced URLs

The **Statistics > Advanced** tab provides a list of *Advanced URLs*. Symantec Technical Support might direct you to these links to provide additional information during troubleshooting.

Chapter 32: Configuring an Application Delivery Network

This section describes how to set up an application delivery network (ADN). It provides basic conceptual and procedural information required to configure ADN. For more detailed information about the recommended ADN deployments, refer to the *Acceleration WebGuide*.

Topics

Refer to the following topics:

- ❑ [Section A: "ADN Overview" on page 712](#)
- ❑ [Section B: "Configuring an ADN" on page 724](#)
- ❑ [Section C: "Securing the ADN" on page 739](#)
- ❑ [Section D: "Configuring Load Balancing" on page 746](#)
- ❑ [Section E: "Configuring Advanced ADN Settings" on page 749](#)
- ❑ [Section F: "Monitoring the ADN" on page 758](#)
- ❑ [Section G: "Related CLI Syntax to Configure an ADN" on page 765](#)
- ❑ [Section H: "Policy" on page 767](#)
- ❑ [Section I: "Troubleshooting" on page 768](#)

Section A: ADN Overview

An Application Delivery Network (ADN) is the core of Symantec's WAN optimization solution. An ADN defines the framework that enables application acceleration between various corporate offices separated by a WAN. In an ADN, ProxySG appliances are integrated into the network to provide visibility, acceleration, and control for traffic sent over the WAN, including:

- ❑ Web (HTTP)
- ❑ Secure Web (SSL)
- ❑ File sharing (CIFS)
- ❑ Microsoft Outlook/Exchange (MAPI)
- ❑ DNS
- ❑ Live and on-demand streaming (Flash RTMP, RTSP, MMS, streaming over HTTP)
- ❑ Other TCP-based applications

With ADN, ProxySG appliances are configured as *ADN nodes*, meaning that they are configured to speak the ADN protocol. When an ADN node intercepts application traffic that has been configured for acceleration, it forms a TCP connection, called a *tunnel*, with the upstream ADN node. The two nodes, called *ADN peers*, send application requests and responses across the tunnel and employ the ADN acceleration techniques that are appropriate for the specific application. The ADN node that intercepts client traffic is referred to as an *ADN Branch peer*; the ADN node that accepts the tunnel connection on the other end of the WAN is called an *ADN Concentrator peer*. An individual ProxySG can act as both an ADN Concentrator peer and an ADN Branch peer; the only difference is its role in a specific tunnel.

The following sections describe the ADN concepts you should understand before configuring ADN:

- ❑ ["ADN Acceleration Techniques"](#) on page 713
- ❑ ["ADN Tunnel Types"](#) on page 714
- ❑ ["ADN Modes"](#) on page 715
- ❑ ["Multiple Concentrators in a Transparent ADN Deployment"](#) on page 716
- ❑ ["ADN Load Balancing"](#) on page 718
- ❑ ["ADN Security"](#) on page 721

ADN Acceleration Techniques

The ProxySG appliances in the ADN apply the following application acceleration techniques appropriate to each application that you want to optimize.

- ❑ **Protocol optimization** — Includes two types of optimizations: Application-layer optimizations and TCP-layer optimizations. Application-layer optimizations improve performance and mitigate the effects of WAN latency, especially for chatty/inefficient protocols like CIFS. Application-layer optimizations include techniques such as read-ahead, pipelining/prefetch, and meta-data caching. TCP-layer optimizations include a variety of techniques to improve link throughput across various WAN environments such as Multi-Protocol Label Switching (MPLS) links, satellite links, or congested/lossy networks.
- ❑ **Object caching** — Reduces latency and bandwidth consumption by caching application data such as CIFS files, Web pages or graphics, and other objects on ProxySG appliances at client sites so that requests are served locally. You can also prepopulate ProxySG appliances with commonly requested content.
- ❑ **Byte caching** — Reduces bandwidth usage by replacing byte sequences in traffic flows with reference tokens. The byte sequences are stored in a byte cache—called a byte-cache dictionary—on a pair of ProxySG appliances at each end of the WAN. When a matching byte sequence is requested again, the ProxySG transmits a token instead of the byte sequence. This acceleration technique is especially beneficial when users make small changes to large documents because the ProxySG only needs to transmit the change across the WAN rather than retransmitting the entire document.

Note: To increase throughput in its tunnels, ADN uses an *adaptive byte caching* mechanism that automatically adjusts byte caching to the amount of disk I/O latency the ProxySG is experiencing. As a ProxySG contends with increasing traffic levels, disk I/O increases and eventually becomes a bottleneck. In these situations, ADN scales back on disk reads and writes of the byte cache. Disk I/O is performed only when it can produce significant byte-caching gain. The end result is higher throughput in ADN tunnels.

Although the byte-cache dictionary is automatically created and sized, there may be times when you must manually modify its size. See "[Configuring the Byte-Cache Dictionary Size](#)" on page 752 for more information.

You can control how long byte-cached data is stored in the dictionary by assigning a *retention priority* to a particular service. If you want to keep certain types of data in the byte cache for as long as possible, set a high retention priority for the service. Or for data that isn't likely to get much benefit from byte caching, you can set a low retention priority for the related service. The default retention priority is "normal." For details on configuring the retention priority, see "[Creating Custom Proxy Services](#)" on page 120.

- ❑ **Compression** — Uses a variety of algorithms to remove extraneous/predictable information from the traffic before it is transmitted. The information is reconstituted at the destination based on the same algorithms. Compression further reduces the size of the content transferred over the network, enabling optimized bandwidth usage and response time to the end user.
- ❑ **Bandwidth management** — Prioritizes and/or limits bandwidth by user or application, allowing WAN usage to reflect business priorities. You can create bandwidth rules using over 500 attributes, such as application, website, URL category, user/group, and time/priority.

ADN Tunnel Types

When an ADN Branch peer intercepts application traffic for optimization, it initiates a TCP connection with the ADN Concentrator peer at the site hosting the application server. This TCP connection between peer ProxySG appliances is called an *ADN tunnel*.

The tunnel type determines the extent to which the packet header information (source IP address, destination IP address, and destination port) from the original packet is retained as the packet travels from client to server across the ADN. There are three types of ADN tunnels as follows:

- ❑ **Transparent** — With a transparent tunnel connection, the original destination IP address and port are maintained. Depending on the desired level of transparency, the connection over the WAN can use the original client's IP address or the IP address of the ADN Branch peer. Transparent tunnels are enabled by default; no additional configuration is required. To use transparent tunnels, the ADN Concentrator peer must be deployed in-path or virtually in-path (and, if you want to use the reflect client IP feature, the ADN Branch peer must be in-path or virtually in-path also). Transparent tunnels are not reused, therefore the ProxySG must use additional resources to create new tunnels.
- ❑ **Translucent** — With a translucent tunnel connection, the ADN Branch peer uses its own address as the source IP address and the ADN Concentrator peer's IP address as the destination IP address while retaining destination port of the server. When you use translucent ADN tunnels, all client traffic is aggregated at the ADN Concentrator peer and you cannot determine traffic use by a specific client, but will be able to see overall traffic by server ports. Use translucent tunnels when the ADN Branch peer is in-path or virtually in-path and the ADN Concentrator peer is out-of-path and there is a need to preserve WAN statistics by service port. For information on creating Translucent tunnels, see ["Enabling Translucent Tunnels"](#) on page 732.

- ❑ **Explicit** — With an explicit tunnel connection, the ADN Branch peer uses its own address as the source IP address and the ADN Concentrator peer's IP address as the destination IP address. Additionally, it uses a destination port number of 3035 (plaintext) or 3037 (secure) by default. Explicit tunnels do not provide granular metrics about which servers and clients use the most network resources. If you are connecting to an ADN Concentrator peer that has been deployed out-of-path, you must use explicit or translucent tunnels. For information on creating Explicit tunnels, see ["Enabling Explicit Tunnels"](#) on page 733.

To establish the tunnel, the ADN Concentrator peer and the ADN Branch Peer must be able to communicate over the tunnel listening port, which is 3035 (plaintext) or 3037 (secure) by default. In an out-of-path deployment, the explicit tunnel and the control connection are established on this port. On an in-path or virtually-in path deployment, the control connection for the transparent or translucent tunnel is established on this port. If the ADN Concentrator peer and the ADN Branch peer cannot communicate over this control connection, byte-cache dictionary synchronization and other non-application-related activities will fail.

Note: ADN devices identify transparent and translucent tunnels by placing a custom TCP option inside the TCP headers. Network devices that remove or modify values found in the fields of the TCP header will cause these tunnels to fail. Intermediary network devices that perform deep packet inspection or NAT firewalls might remove required TCP option information.

ADN Modes

The ADN mode that is configured determines which peers an ADN node can form tunnel connections with. There are two ADN modes as follows:

- ❑ **Open** — An ADN peer is allowed to form a transparent tunnel connection with any other ADN peer.
- ❑ **Closed** — ADN nodes can only establish accelerated tunnel connections with peers in its ADN. In this configuration, you must configure a Primary ADN manager and, optionally, a Backup ADN Manager to manage ADN membership. The ADN manager(s) can be ADN nodes or they can be dedicated ProxySG appliances. In a closed ADN, every ADN peer must connect to the ADN manager(s) in order to become part of the ADN. For instructions on configuring a closed ADN, see ["Configuring a Closed ADN"](#) on page 725.

By default, an ADN operates in Open mode and an ADN manager is not required. This is called *Open-unmanaged* mode (see ["Configuring an Open-unmanaged ADN"](#) on page 724). This allows you to get your ADN up and running quickly and easily. However, because the ADN management functions are not available in an Open-unmanaged ADN, the following are not supported in this configuration:

- ❑ Explicit tunnel connections (including ProxyClient and out-of-path deployments)

- ❑ Load balancing (explicit or transparent)
- ❑ Internet Gateway
- ❑ Manager authorization in secure ADN

To enable any of these services, you must configure an ADN manager and connect the ProxySG appliances that require the services to it. You do not need to connect all ProxySG appliances to the ADN manager. Mixed acceleration networks, in which some open nodes connect to a manager and some do not, is called an *Open-managed ADN* (see "[Configuring an Open-managed ADN](#)" on page 725). The ADN mode is defined on the ADN manager, if there is one. If there is no manager, the ADN mode is Open by default.


Upstream ADN Concentrators

The ADN concentrator is the ProxySG located at the data center in an ADN deployment. For information on how the ADN handles situations where there are multiple concentrators between the client and the server, see "[Multiple Concentrators in a Transparent ADN Deployment](#)" below. For details on how the ADN handle situations where an upstream concentrator is not discovered, see "[Discovery of Upstream Concentrators](#)" on page 718.


Multiple Concentrators in a Transparent ADN Deployment

In transparent ADN deployments where branch office traffic goes through multiple concentrators on its way to and from an origin content server (OCS), you will want to ensure that the ADN tunnel extends across the entire path, allowing the ADN traffic to be optimized from end to end. To achieve this benefit, you enable the *last peer detection* feature on the intermediate concentrators. This feature sends out probes to locate the last qualified peer—the upstream concentrator that has a valid SSL license, closest to the connection's destination address; an ADN tunnel is formed between the branch ProxySG and the last peer en route to the OCS. If there is a concentrator in the path that does not support last peer detection or has it disabled, the transparent tunnel is formed with that concentrator.

Without this feature, the ADN tunnel ends at the first qualified concentrator in the path, as shown in the topology below. The traffic is optimized over this partial segment of the path to the origin content server (OCS). Traffic is not optimized over the rest of the path to the OCS.



Contrast the above illustration with the one shown below. The second illustration shows how the ADN tunnel is lengthened when the last peer detection feature is enabled on the intermediate concentrators. This feature results in the longest ADN tunnel, allowing the traffic to be optimized over the entire path.



Supported ADN Deployments

Last peer detection can be used in transparent ADN deployments including the ones listed below:

- ❑ Physically inline or virtually inline (WCCP) transparent deployments
- ❑ Open ADN mode, managed or unmanaged
- ❑ Closed ADN mode
- ❑ Transparent load balancing deployments
- ❑ Secure ADN
- ❑ Reflect Client IP enabled or disabled on the branch and concentrators
- ❑ SGRP redundancy support on concentrator side

See ["Enabling Last Peer Detection on Transparent Tunnels"](#) on page 733.

Limitations

- ❑ When using last peer detection in a deployment where traffic to an OCS is distributed by a load balancer, there should be a concentrator in each potential path to the OCS. This allows the traffic to be optimized irrespective of the path that the load balancer decides upon.
- ❑ This feature is not operational when the concentrator is performing HTTP proxy processing. For accelerated HTTP traffic, an intermediate concentrator with HTTP proxy processing enabled will not attempt to detect any upstream concentrators and will terminate any inbound transparent tunnels carrying HTTP traffic. Note that the HTTP proxy processing feature has been deprecated.

Discovery of Upstream Concentrators

When an ADN Branch peer intercepts application traffic that has been configured for acceleration, it tries to form a tunnel with the upstream ADN Concentrator peer and the traffic is optimized with ADN acceleration techniques that are appropriate for the specific application. But if an upstream concentrator is not discovered, the tunnel cannot be formed and the traffic cannot be optimized with ADN compression or byte caching. In this situation, you may prefer for the Branch peer to bypass the connection to save memory and CPU resources on the ProxySG. The ProxySG offers a setting that controls whether connections should be bypassed if an upstream concentrator is not detected. This feature is referred to as *bypass-if-no-concentrator*.

Conditions for Activating Bypass Feature

The *bypass-if-no concentrator* feature only applies to services using the TCP Tunnel proxy with ADN active, when an upstream concentrator is not discovered. Several other conditions also apply; see ["Bypass TCP Tunnel Connections When No Concentrator"](#) on page 734.

Supported ADN Deployments

The *bypass-if-no-concentrator* feature can be used in a managed ADN in an explicit deployment as well as in a managed or unmanaged ADN in a transparent deployment. This setting is configured on the ADN Branch peer.

ADN Load Balancing

The way you configure ADN load balancing depends on whether you are using explicit or transparent tunnels. The following sections describe the different types of load balancing:


- ❑ ["Transparent Load Balancing"](#) on page 719
- ❑ ["Explicit Load Balancing"](#) on page 720

Transparent Load Balancing


Configuration of transparent load balancing must be done on each peer in the ADN cluster. Transparent load balancing relies on connection forwarding clusters for proper operation. All peers in an ADN load balancing group must be part of the same connection forwarding cluster. In the context of ADN, connection forwarding relates to how a ProxySG handles the first packet of a request. A decision is made on the first packet about which ADN peer is best to process that request, and subsequently that request is forwarded to that ADN peer from start to finish. If connection forwarding is not set up correctly, load balancing fails. For information on connection forwarding, see ["TCP Connection Forwarding"](#) on page 857. For information on how to configure transparent load balancing, see ["Configuring Transparent Load Balancing"](#) on page 746.

If you are using a transparent deployment, you have two options for load balancing as follows:

- ❑ You can use a ProxySG appliance as a load balancer. In this configuration, the ProxySG that is configured as the load balancer makes the decision about which peer receives which traffic.



- ❑ You can use a WCCP router or L4 switch as an external load balancer. In this configuration, the individual peers in the ADN cluster make the load balancing decision. This configuration is a little more difficult because the WCCP router or L4 switch must be configured on each system in the cluster. In this scenario, the router or switch cannot guarantee ADN peer affinity because the router cannot use the peer ID as input for its hash. Because of this, the ADN peers make the actual routing decisions.



Explicit Load Balancing

If you are using explicit tunnels, you have two options when configuring load balancing:

- ❑ **Server subnet configuration** — In this configuration, you have multiple ProxySG appliances fronting the same IPv4 and/or IPv6 server subnets. Using a hashing function, each ProxySG determines its preferred peer to which it will route traffic destined for the load-balanced subnet. In this configuration, no allowance is made for equalizing load among different sized hardware in the same ADN cluster.
- ❑ **External load balancer configuration** — In this configuration, you configure an external load balancer to front a group of ADN peers. The external load balancer distributes the load among the peers it fronts using client/IP address affinity.

For more information, see ["Configuring Explicit Load Balancing"](#) on page 747.

ADN Security

The choices for securing your ADN depend on the ADN mode you are using. Many of the ADN security features rely on the ADN manager for enforcement ("[Managed ADN Security](#)" on page 721); therefore if your ADN is operating without a manager ("[Unmanaged ADN Security](#)" on page 721), you will not be able to use all of the security features. By default, none of the ADN security features are enabled.

Unmanaged ADN Security

If your ADN is operating in Open-unmanaged mode, any ADN node can form transparent tunnel connections with any other ADN node. Thus, your ADN nodes are at risk for attack from systems outside your network.

To ensure that your ADN nodes only connect to authorized ADN nodes, you must deploy your own public key infrastructure (PKI) within your ADN and then secure the tunnel connections the ADN peers use. By issuing certificates to authorized ADN nodes only, you ensure that your ADN nodes will only be able to form tunnel connections with other authorized ADN nodes. For more information on securing an Open-unmanaged ADN, see "[Securing an Unmanaged ADN](#)" on page 739.

Managed ADN Security

If you are using an ADN manager, you can use the secure ADN features. Secure ADN requires an appliance certificate for each ADN peer—including the ADN manager and backup manager—for identification. You can provide your own device appliance certificates or obtain Blue Coat-issued appliance certificates from the Blue Coat CA server. To enable secure ADN, you must enable the appliance authentication profile for the ADN to use before configuring any other security parameters. Secure ADN provides the following features:

- ❑ "[ADN Peer Authentication](#)"
- ❑ "[ADN Peer Authorization](#)"
- ❑ "[ADN Connection Security](#)"

For more information, see "[Securing a Managed ADN](#)" on page 740.

ADN Peer Authentication

In secure ADN mode, full mutual authentication can be supported between the ADN manager and the nodes that are connected to it and between ADN peers. To use authentication, each node must have an SSL certificate and have an SSL device profile configured. For more information on managing appliance certificates, see "[Authenticating a ProxySG](#)" on page 1289. For information on enabling device authentication on your ADN nodes, see "[Enabling Device Authentication](#)" on page 740.

ADN Peer Authorization

If authorization is enabled, the ADN manager must authorize a node before it is allowed to join the ADN as follows:

- ❑ When an ADN peer comes up, it contacts the ADN manager for routing information.
- ❑ The ADN manager extracts the device ID from the connecting ADN peer's appliance certificate and looks for the device ID in its approved list of ADN peers.
 - If the device is on the approved list, a `REQUEST-APPROVED` response is sent, followed by the route information, and the peer joins the network.
 - If the **Pending Peers** option is enabled and the device is not on the approved list, the ADN manager adds the connecting peer's device ID to a pending-peers list and sends a `REQUEST-PENDING` response. After the peer is moved to the **Approved** list by the administrator, a `REQUEST-APPROVED` response is sent, followed by the route information, and the peer joins the network.
 - If the **Pending Peers** option is not enabled and a peer is not on the approved list, the ADN manager sends a `REQUEST-DENIED` response and closes the connection. The connecting peer closes the connection and updates its connection status.
 - If a peer is deleted from the approved list, the ADN manager broadcasts a `REJECT-PEER` to all peers to delete this peer and terminate any existing ADN connections to it. No new connections are routed through the deleted ADN peer.

For information on configuring authentication and authorization on each ADN peer, see ["Securing a Managed ADN"](#) on page 740.

ADN Connection Security

By default, ADN routing and tunnel connection requests are unauthenticated and all ADN protocol messaging and compressed application data are transferred in plaintext. For maximum security, you can configure the ADN to secure ADN routing and tunnel connections using standard SSL protocol, which provides authentication, message privacy, and message authenticity security services, regardless of the application traffic that is being accelerated or tunneled.

In secure ADN mode, you can specify that the ADN manager and tunnel use secure mode to listen for routing and tunnel requests. When ADN connection security is enabled, any existing plain outbound connections are dynamically secured by activating SSL according to the `secure-outbound` setting.

The following table describes secure outbound behavior with various applications.

Table 32–1 Secure Outbound Behavior

Secure-Outbound Setting	Routing Connections	Application Connections		
		CIFS	SSL Proxy Intercept Mode	SSL Proxy Tunnel Mode
None	Plain Text	Plain Text	Bypass ADN	Bypass ADN
Secure Proxies	Encrypted	Plain Text	Encrypted	Encrypted by application
All	Encrypted	Encrypted	Encrypted	Encrypted by application

For information on optimizing and securing ADN tunnels, see "[Securing the ADN](#)" on page 739 and "[Configuring Advanced ADN Settings](#)" on page 749.

Section B: Configuring an ADN

This section discusses the following topics:

- ❑ ["Introduction to Configuring an ADN"](#)
- ❑ ["Configuring an Open-unmanaged ADN" on page 724](#)
- ❑ ["Configuring an Open-managed ADN" on page 725](#)
- ❑ ["Configuring a Closed ADN" on page 725](#)
- ❑ ["Switching ADN Modes" on page 728](#)
- ❑ ["Enabling Explicit ADN Connections" on page 729](#)
- ❑ ["Configuring IP Address Reflection" on page 735](#)

Introduction to Configuring an ADN

The steps that are required to set up an ADN depend on the ADN mode you plan to use and the type of tunnels (explicit or transparent) that you are using as follows:

- ❑ ["Configuring an Open-unmanaged ADN" on page 724](#)
- ❑ ["Configuring an Open-managed ADN" on page 725](#)
- ❑ ["Configuring a Closed ADN" on page 725](#)
- ❑ ["Switching ADN Modes" on page 728](#)
- ❑ ["Enabling Explicit ADN Connections" on page 729](#)
- ❑ ["Configuring IP Address Reflection" on page 735](#)
- ❑ ["Enabling ProxyClient Support" on page 737](#)

Note: In addition to the tasks you must perform on the ProxySG appliance to enable acceleration, you must also make sure that your firewall is configured to allow tunnel connections between your ADN Concentrator peers and your ADN Branch peers for all deployment types (in-path, virtually in-path, or out-of-path; Open and Closed). To do this, open the tunnel listening port on the ADN Concentrator side of the firewall. By default, this port is set to 3035 (plain) and 3037 (secure). This port is used to create the control connection for the tunnel, which is used to synchronize ADN byte-cache dictionaries and other non-application-related activities. In explicit deployments, this port is also required to establish the explicit tunnel.

Configuring an Open-unmanaged ADN

An Open-unmanaged ADN is an ADN in which any ADN node can connect transparently to any other ADN node. There is no ADN manager and all nodes must be using transparent tunnels (and therefore must be deployed in-path or virtually in-path).

Open-unmanaged ADN is the default and it requires very little configuration. To set up an ADN in Open-unmanaged mode, complete the following steps:

- ❑ Install the ProxySG appliances that will be your ADN nodes in-path or virtually in-path. For instructions, refer to the *Quick Start Guide* for the specific ProxySG platform.
- ❑ To accelerate applications other than those that are accelerated by default, configure the corresponding Proxy Services. For information on configuring proxy services, see ["Configuring a Service to Intercept Traffic"](#) on page 117.
- ❑ If you did not enable acceleration during setup, you must enable it as follows.

To enable ADN on an Open-unmanaged ADN node:

1. Select **Configuration > ADN > General**.
2. Select **Enable Application Delivery Network**.
3. Verify that the **Primary ADN Manager** and **Backup ADN Manager** are set to **None**.
4. Click **Apply**.
5. Repeat these steps on each ADN node.

Configuring an Open-managed ADN

In an Open-managed ADN, any ADN node can connect transparently to any other ADN node. However, some ADN nodes are also configured to use an ADN manager. This is a common deployment when you have some sites that require services that rely on an ADN manager, such as ProxyClient, but you still want your ADN to operate in Open mode.

To operate in Open-managed mode:

1. Configure a Primary ADN manager and optionally a Backup ADN manager. See ["Configuring the ADN Managers and Enabling ADN"](#) on page 726 for instructions.
2. For each ADN node that needs to be managed, configure the node to connect to the ADN manager(s). See ["Configuring the ADN Managers and Enabling ADN"](#) on page 726 for instructions.
3. Nodes that do not need to be managed (that is, they do not require any ADN manager services) do not need any additional configuration. You can configure them as described in ["Configuring an Open-unmanaged ADN"](#) on page 724.

Configuring a Closed ADN

In a Closed ADN, an ADN node is only allowed to connect to peers that are in their ADN, as defined by an ADN manager. Therefore, to configure Closed ADN define your ADN manager(s) and configure every ADN node to connect to the manager(s). An ADN manager can be any ADN node or it can be a dedicated ProxySG appliance (recommended in large deployments).

To configure a Closed ADN you must:

- ❑ Configure a Primary ADN manager and enable ADN on it. See ["Configuring the ADN Managers and Enabling ADN"](#) on page 726.
- ❑ (Optional) Configure a Backup ADN manager and enable ADN on it. Configuring a Backup ADN manager is recommended, but not required. See ["Configuring the ADN Managers and Enabling ADN"](#) on page 726.
- ❑ Configure each ADN node to connect to the ADN manager(s). See ["Configuring the ADN Managers and Enabling ADN"](#) on page 726.
- ❑ Configure the Primary ADN manager and the Backup ADN Manager (if one exists) to operate in Closed mode. See ["Setting the ADN Mode"](#) on page 728 for instructions.

Configuring the ADN Managers and Enabling ADN

If you plan to run your ADN in Closed mode or Open-managed mode, you must configure a Primary ADN manager and optionally a Backup ADN manager. Begin by configuring the ProxySG appliance(s) that will function as the Primary and Backup ADN managers. After you configure the managers, you must configure each ADN node that you want to be managed to connect to the manager(s).

Note: When upgrading managed ADN deployments to a release that supports IPv6 on ADN (SGOS 6.2.4 or higher), the ProxySG that is functioning as the ADN manager must be upgraded before the managed nodes. The manager should continue to be assigned a reachable IPv4 address until all managed nodes have been upgraded. A managed node that has been upgraded to a release that supports IPv6 on ADN (SGOS 6.2.4 or higher) can use either IPv4 or IPv6 to connect to the previously upgraded manager.

To define the ADN Managers and enable ADN on each node:

1. Select **Configuration > ADN > General**.

2. **Primary ADN Manager:**

- If this ProxySG is the Primary ADN manager, select **Self**.
- For other ADN nodes, select **IP Address** and enter the IPv4 or IPv6 address of the ProxySG that is configured as the Primary ADN manager. A Primary ADN manager is required if you are in Open-managed mode or Closed mode.
- If there is no ADN manager (Open-unmanaged mode only), select **None**.

3. **Backup ADN Manager:**

- If this ProxySG is the Backup ADN manager, select **Self**.
- For other ADN nodes, select **IP Address** and enter the IPv4 or IPv6 address of the ProxySG that is configured as the Backup ADN manager (if any). A Backup ADN manager is recommended, but only required if you have ADN nodes deployed out-of-path. The Backup ADN Manager does not need to be the same IP version as the Primary ADN Manager.
- If you do not have a Backup ADN manager, select **None**.

4. **Manager Ports:** The ports are set to 3034 (for plain routing connections) and port 3036 (for secure routing connections) by default. However, to reduce the number of ports that you use for ADN, you can change the manager ports to the same port numbers used for ADN tunnel connections. By default, ADN tunnel connections use ports 3035 (plain) and 3037 (secure), however, you can change these values.

5. Select **Enable Application Delivery Network**.

6. Click **Apply**.


Setting the ADN Mode

The ADN mode determines what peers an ADN node can connect to. In Open mode, the default, an ADN node can connect to any other ADN node. In Closed mode, ADN nodes cannot connect unless they are in the same ADN as defined by an ADN manager. You define the mode by toggling a option on the **Peer Authorization** tab on the ADN manager(s).

To switch the mode, you must perform this procedure on both the Primary ADN manager and the Backup ADN manager (if there is one).

To set the ADN mode:

1. Select **Configuration > ADN > Manager > Peer Authorization**.



2. Set the **Allow transparent tunnels only within this managed network** option as follows:
 - To set the ADN mode to Closed, make sure the option is checked.
 - To set the ADN mode to Open, make sure the option is cleared.
3. Click **Apply**.

Switching ADN Modes

When switching from one ADN mode to another, you must consider the order in which you transition each node as described in the following sections:

- ❑ "Switching from a Closed ADN to an Open ADN" on page 728
- ❑ "Switching from an Open ADN to a Closed ADN" on page 729

Switching from a Closed ADN to an Open ADN

To switch from a Closed ADN to an Open ADN:

To switch the mode from Closed to Open, you simply uncheck the **Allow transparent tunnels only within this managed network** option on the ADN manager(s) as described in "To set the ADN mode:" on page 728.

Switching from an Open ADN to a Closed ADN

To switch from an Open ADN to a Closed ADN:

1. Configure an ADN manager, if one is not already enabled. See ["Configuring the ADN Managers and Enabling ADN"](#) on page 726.
2. Configure each ADN node that you want to be part of the Closed ADN to connect to the ADN manager(s). See ["Configuring the ADN Managers and Enabling ADN"](#) on page 726.
3. If any of the nodes need to advertise server subnets, set up the advertisements. See ["Advertising Server Subnets"](#) on page 729.
4. After you configure the ADN manager(s) and connect each node to them, change the ADN mode to Closed as described in ["Setting the ADN Mode"](#) on page 728.

Enabling Explicit ADN Connections

If any of your ADN nodes is deployed out-of-path or if you plan to use explicit or translucent tunnels, you will need to perform some additional configuration steps as follows:

- ❑ If any of your ADN nodes is deployed out-of-path, you must advertise the subnets it serves. See ["Advertising Server Subnets"](#) on page 729.
- ❑ Transparent tunnels are created automatically. However, if an ADN Branch peer receives explicit routes from an ADN Concentrator peer, the type of tunnel that the ADN Branch peer will form with the ADN Concentrator peer depends on the tunnel mode settings. If the ADN Branch peer is allowed to form transparent tunnels and the ADN Concentrator is configured to prefer transparent tunnels, the ADN Branch peer will form a transparent tunnel. If it cannot form a transparent tunnel, it will check to see if the ADN Concentrator peer is configured to preserve the destination port; if so, it will form a translucent tunnel. Otherwise it will form an explicit tunnel. See ["Configuring the Tunnel Mode"](#) on page 730.


Advertising Server Subnets

If you deploy an ADN Concentrator peer out-of-path, you must advertise the subnets to which it is connected so that the ADN Branch peers can establish connections with it.

Note: You can also configure the exempt subnet capability through policy that allows you to disable ADN tunnels for specific connections. For more information, refer to the *Content Policy Language Guide*.

To advertise server subnets for this peer:

1. Select **Configuration > ADN > Routing > Server Subnets**.



2. To add a subnet, click **Add**. The Add IP/Subnet dialog box is displayed.
3. Define a subnet as follows and then click **OK**:
 - **IP address**: Enter an IPv4 or IPv6 address.
 - **Prefix length or subnet mask**: Specify the prefix length (for IPv6) or subnet mask (for IPv4).
4. Repeat steps 2 and 3 for each subnet.
5. To remove subnets, do one of the following:
 - To remove an individual subnet, select the subnet and click **Remove**.
 - To remove all subnets, click **Clear all**.
6. Click **Apply**.

Configuring the Tunnel Mode

An ADN tunnel is a TCP connection established between an ADN Branch peer and an ADN Concentrator peer and is used to optimize inbound and outbound traffic. ADN tunnels are of three types: Transparent, Translucent, and Explicit.

Transparent tunnels can be used when the ADN Concentrator peer is deployed in-path or virtually in-path. They are enabled by default and require no additional configuration. However, transparent and translucent tunnel connections require a control connection, which is used to synchronize ADN byte-cache dictionaries and other non-application-related activities. This requires that you open the tunnel listening port (3035/3037 for plain/secure connections by default) on the ADN Concentrator side of the connection to ensure successful acceleration over the tunnel.

Note: Starting in SGOS 5.5, the ADN protocol speeds up transparent tunnel establishment. To take advantage of this feature, the Branch peers must be running this version or higher. Make sure to upgrade ADN nodes in the following order: First, upgrade the Primary ADN manager and Backup ADN manager; next, upgrade all appliances that only act as Concentrator peers; finally, upgrade all appliances that act as Branch peers.

Note that a Concentrator peer will intercept a transparent tunnel from a Branch peer only when it is configured with at least one address of the same address family (IPv4/IPv6) as the destination (OCS) address.

If you have an out-of-path ADN Concentrator peer, you must use explicit tunnels or translucent tunnels. If an ADN Branch peer receives advertised explicit routes from an ADN Concentrator peer, it must determine what type of tunnel to establish based on the tunnel mode settings. If the routing preference on the ADN Concentrator peer is set to prefer transparent tunnels, the ADN Branch peer attempts to create a transparent tunnel if it is allowed to. If not, it checks whether the ADN Concentrator peer is configured to preserve the destination port, and, if so it will attempt to establish a translucent tunnel. Otherwise, it establishes an explicit tunnel. For information on each type of tunnel and when to use it, see ["ADN Tunnel Types"](#) on page 714.

The following sections describe how to configure the settings that are used to configure the tunnel mode:

- ❑ ["Setting the Routing Preference"](#) on page 732
- ❑ ["Enabling Translucent Tunnels"](#) on page 732
- ❑ ["Enabling Explicit Tunnels"](#) on page 733
- ❑ ["Enabling Last Peer Detection on Transparent Tunnels"](#) on page 733
- ❑ ["Bypass TCP Tunnel Connections When No Concentrator"](#) on page 734


Note: The proxy processing feature has been deprecated. The Proxy Processing tab has been removed from the Management Console, but the feature can still be configured via the CLI. Since proxy processing will be completely removed from an SGOS release in the near future, Symantec recommends that you discontinue using this feature and deploy a separate secure web gateway to handle proxy processing.

Setting the Routing Preference

If your ADN has a mixed tunnel environment (some explicit tunnels and some transparent tunnels), you can specify what type of tunnels you would prefer the ADN node to use. You can configure the ADN node so that transparent tunnels are used whenever possible.

To configure the ADN node to prefer transparent tunnels:

1. Select **Configuration > ADN > Routing > Advanced**.




2. Select the **Tell ADN peers to prefer transparent connections over advertised routes** option.
3. Click **Apply**.

Enabling Translucent Tunnels

If you are using explicit tunnels, but you would prefer to preserve the destination TCP port number, you can configure the translucent tunnel mode as follows.

To enable translucent tunnels:

1. Select **Configuration > ADN > Tunneling > Connection**.



2. Select the **When a route is available, preserve the destination TCP port number when connecting to the ADN peer**.
3. Click **Apply**.

Enabling Explicit Tunnels

If you are using explicit tunnels, you enable them as follows:

To enable explicit tunnels:

1. Select **Configuration > ADN > Tunneling**.
2. Clear the **Connect using ADN transparent tunneling when possible** option.
3. Click **Apply**.

Enabling Last Peer Detection on Transparent Tunnels

When your transparent ADN deployment has multiple concentrators between the branch office and the OCS, you should enable the last peer detection feature. For details on this feature, see ["Multiple Concentrators in a Transparent ADN Deployment"](#) on page 716.

To enable last peer detection:

The last peer detection feature is automatically enabled on fresh installations but is disabled on upgraded systems. Although it doesn't hurt to enable the feature on every ProxySG on the path, it is only required to be enabled on the intermediate concentrators. If you want a particular concentrator to terminate the transparent tunnel, you can disable last peer detection on that ProxySG.

1. Select **Configuration > ADN > Tunneling > Connection**.

2. Select **Automatically detect last ADN peer on path to OCS**.
3. Click **Apply**.

When the intermediate concentrators are configured for last peer detection, you would expect that the active sessions on these ProxySG appliances would show the connections being bypassed due to "upstream ADN peer detection." On the other hand, the active sessions on the last concentrator would show the connections from the branch office being intercepted and optimized. See ["Reviewing ADN Active Sessions"](#) on page 759.

Bypass TCP Tunnel Connections When No Concentrator

For a TCP Tunnel connection with no special policy controls, if the underlying ADN optimization can not be provided because of the absence of an upstream ADN concentrator, there is little value in intercepting the connection. In such scenarios, to save memory and CPU resources on the ProxySG, you can configure the Branch peer to bypass the TCP Tunnel connection if an upstream ADN concentrator is not discovered.

For additional information about the bypass-if-no-concentrator feature, see ["Discovery of Upstream Concentrators"](#) on page 718.

On fresh installations of the Acceleration Solution, the bypass-if-no-concentrator feature is automatically enabled, with the **Only when no applicable policy exists** option selected. Note that bypass-if-no-concentrator is disabled by default on systems upgraded from a pre-6.3 software version.

To configure bypass-if-no-concentrator:

1. Select **Configuration > ADN > Tunneling > Connection**.

The screenshot shows the 'Connection' tab in the ProxySG configuration interface. Under the 'Inbound' section, 'Plain Tunnel Port' is set to 3035 and 'Secure Tunnel Port' is set to 3037. Under the 'Outbound' section, the following options are checked:

- ☒ Connect using ADN transparent tunneling when possible
- ☒ When a route is available, preserve the destination TCP port number when connecting to the ADN peer
- ☒ Automatically detect last ADN peer on path to OCS
- ☒ When no concentrator is found, bypass TCP Tunnel traffic
 - ☐ Whenever possible
 - ☒ Only when no applicable policy exists

Arrows labeled 2 and 3 point to the 'When no concentrator is found, bypass TCP Tunnel traffic' option and its sub-options respectively.

2. To disable this feature, clear the **When no concentrator is found, bypass TCP Tunnel traffic** option.
3. To enable bypass-if-no-concentrator, select **When no concentrator is found, bypass TCP Tunnel traffic** and then choose one of the following:

Whenever possible—With this option, all the following conditions must be true in order for a connection to be bypassed:

- ADN optimization is enabled on this ProxySG.
- The service has ADN enabled, uses the TCP Tunnel proxy, and is marked for interception.
- Initial policy check allows the connection to go through to the server.
- Early Intercept is disabled.
- Detect Protocol is disabled.
- An upstream Concentrator is not discovered.

Only when no applicable policy exists—This option applies additional conditions that must be met in order for the connection to be bypassed:

- DSCP is set to *preserve* for outbound client/server traffic.
- No bandwidth management class is set for client and server, inbound and outbound flows.
- No forwarding rule applies that redirects the connection to a different upstream server or to a SOCKS proxy
- No URL-rewrite rules apply that affect the server URL setting.

4. Click **Apply**.

When this feature is enabled and a connection is bypassed because an upstream concentrator was not found, the Active Sessions report indicates the reason the connection was bypassed; in the Bypassed Connections tab, the **Details** column lists **No ADN concentrator was discovered** for each bypassed connection.

Client IP	Server	Duration	Bypassed Reason	Service Name	Details
10.9.45.192:3844	10.9.45.198:1352	5 sec	50,471	Lotus Not ...	No ADN concentrator was discovered
10.9.45.192:48603	10.9.45.198:8082	2.8 min	1,794	Blue Coat ...	
10.9.45.193:1027	10.9.45.198:3034	4.2 min	1,110	Blue Coat ...	
10.9.45.193:1048	10.9.45.198:3035	3.2 min	2,023	Blue Coat ...	

Configuring IP Address Reflection

By default, an ADN Branch peer uses its own IP address when creating an ADN tunnel connection with an ADN Concentrator peer. However, in some deployments you can configure the ADN so that the client IP address is retained. This process is called *client IP address reflection*.

Symantec recommends configuring client IP address reflection whenever possible because it provides maximum visibility for network usage statistics and enables user-based access control to network resources.

SGOS 6.2 or later offers independent controls for configuring how the Concentrator peer handles client IP reflection requests from ProxySG peers versus ProxyClient peers. For example, you can have the Concentrator reject client IP reflection requests from ProxyClient peers but allow them from ProxySG peers. In previous releases, when the Concentrator was configured to *deny* reflect client IP requests from branch peers, there was a special hard-coded override that always used the Concentrator's local IP address for ProxyClient tunnel connections; if reflect client IP was set to *allow*, then the client IP would be reflected.

The way you configure client IP address reflection depends on whether the appliance will act as an ADN Branch peer, an ADN Concentrator peer, or both. Use the following procedures to configure client IP address reflection.

- ❑ If the ADN node will act as an ADN Concentrator peer, see "[To configure client IP address reflection on an ADN Concentrator Peer:](#)" on page 736.
- ❑ If the ADN node will act as an ADN Branch peer, see "[To configure client IP address reflection on an ADN Branch peer:](#)" on page 737.

To configure client IP address reflection on an ADN Concentrator Peer:

1. Select **Configuration > ADN > Tunneling > Network**.

The screenshot shows the 'Network' tab in the 'ADN Tunneling' configuration section. It contains two main settings:

- Reflect Client IP:** A section with the text 'When an inbound ADN tunnel connection requests Client IP reflection:' and three radio button options:
 - ☐ Reject the request
 - ☐ Allow the request and reflect the client IP
 - ☒ Allow the request but connect using a local IP
- TCP Settings:** A section with the text 'ADN Tunnel TCP Window Size:' and two options:
 - ☒ Automatically adjusted
 - ☐ Manual override:

2. Determine the behavior of the ADN Concentrator peer when a ProxySG Branch peer requests client IP reflection for an inbound tunnel connection. The ADN Concentrator peer client IP reflection configuration determines what IP address the ADN Concentrator peer advertises to the origin content server (OCS) as the source address: its own address (referred to as *use local IP*) or the client's IP address (referred to as *reflect client IP*).

The option you select depends mainly on whether or not the ADN Concentrator peer is deployed in-path or virtually in-path between the ADN Branch peer and the OCS, as follows:

- **Reject the request**

Select this option to reject requests to reflect the client IP; as a result, the connection to the ADN Concentrator peer is rejected.

- **Allow the request and reflect the client IP**

Choose this option if the ADN Concentrator peer is deployed in-path or virtually-in path between the ADN Branch peer and the OCS. This option indicates that the return packets will have the client's IP address as the destination address and must be routed back through the same ADN Concentrator peer.

- **Allow the request but connect using a local IP**


Choose this option if the ADN Concentrator peer is deployed out-of-path with respect to the ADN Branch peer and the OCS or if there are asymmetric routing issues in which a server response may not always flow through the ADN Concentrator peer.

Note: You can also modify the TCP window size from this tab. For more information, see ["Modifying the TCP Window Size"](#) on page 751.

3. Determine the behavior of the ADN Concentrator peer when a *ProxyClient* Branch peer requests client IP reflection for an inbound tunnel connection.
4. Click **Apply**.

To configure client IP address reflection on an ADN Branch peer:

1. Select **Configuration > Proxy Settings > General**.



The screenshot shows a configuration window with a 'General' tab selected. There are two expandable sections. The first section, 'Tunnel on Protocol Error', contains a checkbox labeled 'TCP tunnel requests when a protocol error is detected' which is currently unchecked. The second section, 'Reflect Client IP', contains a checkbox labeled 'Reflect client's source IP when connecting to servers' which is currently checked.

2. Select **Reflect client's source IP when connecting to servers**.
3. Click **Apply**.

Enabling ProxyClient Support

The ProxyClient does not advertise routes; instead, it gets routes from the ADN manager. To use the ProxyClient in your ADN, all ADN Concentrator peers that front servers that will accelerate traffic for ProxyClients must specify a Primary ADN manager and optionally a Backup ADN manager.

In other words, to use the ProxyClient in your ADN, you must use either Open-managed or Closed ADN.

To enable ProxyClient support:

1. Configure an ADN manager, if one is not already enabled. See ["Configuring the ADN Managers and Enabling ADN"](#) on page 726.
2. On the ADN manager(s), set the **Manager Listening Mode** to **Plain read-only** (recommended), **Plain-only**, or **Both** as discussed in ["Configuring Connection Security"](#) on page 742.
3. Set the ADN manager(s) tunnel listening mode to **Plain Only** or **Both** (recommended) as discussed in ["Configuring Connection Security"](#) on page 742.
4. On each ADN Concentrator peer fronting servers that the ProxyClients need access to:
 - Configure the ADN Concentrator peer to connect to the ADN manager(s). See ["Configuring the ADN Managers and Enabling ADN"](#) on page 726.

- Advertise the subnets that the ADN Concentrator peer services. See ["Advertising Server Subnets"](#) on page 729.

For more information about setting ADN options for use with ProxyClient, refer to the *ProxyClient Configuration and Deployment Guide*.

Section C: Securing the ADN

The options that are available for securing your ADN depend on whether the ADN is unmanaged (Open-unmanaged mode) or managed (Open-managed or Closed mode).

- ❑ **Secure Unmanaged ADN** — If you are operating in Open-unmanaged mode you cannot use the security features provided by the ADN manager. Additionally, in this mode any ADN node can form a transparent connection with any other ADN node. To ensure that your ADN nodes only connect to authorized ADN nodes, you must deploy your own public key infrastructure (PKI) within your ADN and then secure the tunnel connections the ADN peers use. See ["Securing an Unmanaged ADN"](#) on page 739 for more information.
- ❑ **Secure Managed ADN** — If you are operating in Open-managed or Closed mode, you can use the secure ADN features provided by the ADN manager (including device authentication and authorization and secure routing connections). If you are in Open-managed mode, only managed nodes (nodes that are configured to connect to an ADN manager) can use the secure ADN features. See ["Securing a Managed ADN"](#) on page 740 for more information.

Securing an Unmanaged ADN

To prevent an ADN node in an Open-unmanaged ADN from forming connections with any other ADN node, you can enable an SSL device profile so that the devices must authenticate before forming tunnel connections. Because the default SSL device profile will be the same for all ProxySG appliances, you will need to issue your own certificates and create a new device profile in order for the authentication to be secure.

To secure an unmanaged ADN:

1. Using your own PKI system, generate a certificate for each ProxySG appliance and install them on each appliance along with the certificate for your CA. See ["Manually Obtaining an Appliance Certificate"](#) on page 1292 for instructions on how to import a certificate.
2. Create a CA Certificate List (CCL) for your CA. For information on creating a CCL, see ["Managing CA Certificate Lists"](#) on page 1143.
3. On each ProxySG, create an SSL device authentication profile that references the new certificate keyring. See ["Creating an SSL Device Profile for Device Authentication"](#) on page 1296 for instructions.
4. Enable the new SSL device profile by selecting **Configuration > ADN > General > Device Security**, selecting the **SSL Device Profile** from the drop-down list, and then clicking **Apply**.

5. Configure each ADN node to form tunnels over secure connections only by selecting **Configuration > ADN > General > Connection Security** and then selecting the **Secure Only** option in the **Tunnel Listening Mode** section of the screen. Click **Apply**.

Securing a Managed ADN

If your ADN uses an ADN manager, you can use the following secure ADN features to secure your ADN.

- ❑ **Device authentication** — With device authentication, the ADN manager verifies the node's peer ID before allowing a connection. See "[Enabling Device Authentication](#)" on page 740.
- ❑ **Connection security** — Allows you to secure tunnel and routing connections. See "[Configuring Connection Security](#)" on page 742.
- ❑ **Device authorization** — With device authorization, the ADN manager must approve all peer connections. See "[Enabling Device Authorization](#)" on page 743.

For maximum security, configure the ADN for both device authentication and device authorization. You must configure device authentication before you can configure connection security and device authorization.

Note: Secure tunnel connections for applications such as CIFS, MAPI, TCP Tunnel, HTTP, or HTTPS/SSL, are dependent upon an SSL license.

Enabling Device Authentication


When you configure device authentication, you select an SSL device profile to use to secure your ADN nodes. After you have selected an SSL device profile, the ADN manager will automatically verify a ProxySG appliance's peer ID before allowing it to join the ADN.

You can use the default SSL certificate and default SSL device profile (bluecoat-appliance-certificate) or you can import your own certificates and define a new SSL device profile. For more information on device authentication, see "[Authenticating a ProxySG](#)" on page 1289.

Note: If the device being configured for authentication has Internet access, acquisition of the ProxySG appliance certificate is automatic. If you use your own appliance certificates and profile, or if the affected device does not have Internet access, manual device authentication is required.

To enable device authentication:

1. On each peer, configure a Primary ADN manager and optionally a Backup ADN manager if you haven't already done so. See ["Configuring the ADN Managers and Enabling ADN"](#) on page 726.
2. Select **Configuration > ADN > General > Device Security**.



3. Configure the **Device Security** options:
 - a. **SSL Device Profile:** From the drop-down list, select the device profile you want to use. You can use the default bluecoat-appliance-certificate profile or a custom profile. You must use the same profile on each node in the ADN.
 - b. **Extracted Device ID:** The device ID that was extracted based on the selected profile is automatically displayed.

Note: The device ID is only used for security. The peer ID is the serial number.

- c. To enable authorization, select the **Validate ADN Peer Device IDs** option.
 - If the primary or backup ADN manager is **Self**, you do not need to retrieve the device ID.
 - If the primary or backup ADN manager is a different system, click the **Retrieve Manager IDs** button to retrieve the device ID. Click **Accept** to add the manager device ID to the **Primary Manager Device ID** or **Backup Manager Device ID** field.
4. Click **Apply**.

Configuring Connection Security

By default, ADN routing and tunnel connection requests are unauthenticated and all ADN protocol messaging and compressed application data are transferred in plain text. After you configure a device authentication profile ("[Enabling Device Authentication](#)" on page 740), you can configure connection security as follows.

To configure connection security and define the manager and tunnel listening ports:

1. Select **Configuration > ADN > General > Connection Security**.

The screenshot shows the 'Connection Security' configuration page. It has three tabs: 'General', 'Device Security', and 'Connection Security'. The 'Connection Security' tab is selected. There are three main configuration sections, each with a title and a set of radio buttons:

- Manager Listening Mode** (indicated by arrow 2):
 - ☐ Secure Only
 - ☐ Plain Read-Only
 - ☐ Plain Only
 - ☒ Both
- Tunneling Listening Mode** (indicated by arrow 3):
 - ☐ Secure Only
 - ☐ Plain Only
 - ☒ Both
- Secure-Outbound Mode** (indicated by arrow 4):
 - ☐ Do not secure ADN connections
 - ☒ Secure ADN routing connections and tunnel connections made by secure proxies
 - ☐ Secure all ADN routing and tunnel connections

2. Select a manager listening mode. By default, the ADN manager(s) will listen for requests on both the plain port and the secure port (**Both**) if you have selected a device authentication profile. You can change the manager listening mode by selecting one of the following:
 - **Secure Only** — The ADN manager(s) will listen for requests on the secure port only.
 - **Plain Read-Only** — This mode is recommended if ProxyClient is deployed in your ADN. Currently, ProxyClient does not support secure ADN. For information about using the other modes with the ProxyClient, refer to the *ProxyClient Configuration and Deployment Guide*.
 - **Plain Only** — The ADN manager(s) will listen for requests on the plain port only.

3. Select a tunnel listening mode. By default, the tunnel listening mode will be set to listen for requests on both the plain port and the secure port (**Both**) if you have selected a device authentication profile. You can change the tunnel listening mode by selecting one of the following:
 - **Secure Only** — The tunnel listener will listen for requests on the secure port only. Do not use this mode if you have ProxyClients deployed in your ADN.
 - **Plain Only** — The tunnel listener will listen for requests on the plain port only.
 4. Select a secure-outbound mode. By default, the ProxySG is configured to **Secure ADN routing connections and tunnel connections made by secure proxies**. You can change the secure-outbound mode by selecting one of the following options:
 - **Do not secure ADN connections** — Neither routing nor tunnel connections are secured. Secure proxy connections bypass ADN and go directly to the OCS.
 - **Secure all ADN routing and tunnel connections** — All outbound routing and tunnel connections are secured. Only use this option if the ProxySG platform has capacity to handle the extra overhead.
-
- Note:** You must have an SSL license in order to secure outbound tunnel connections.
-
5. To change the manager listening ports, select **Configuration > ADN > General > General**. The default plain port is 3034; the default secure port is 3036. To consolidate the number for ports required for ADN, you can set the manager listening ports to the same port numbers you use for ADN tunnel connections: 3035 (plain) and 3037 (secure) by default.
 6. To change tunnel listening ports, select **Configuration > ADN > Tunneling > Connection**. The default is plain port is 3035; the default secure port is 3037.
 7. Click **Apply**.

Enabling Device Authorization

With device authorization, a ProxySG will not be allowed to join the ADN until it has been approved by the Primary ADN manager and Backup ADN manager (if configured). You must enable authentication on all ADN nodes before you can enable authorization. For instructions on enabling authentication, see ["Enabling Device Authentication"](#) on page 740.


This section discusses the following topics:

- ["Managing Authorized Peers"](#) on page 744
- ["Approving a Peer"](#) on page 745

Managing Authorized Peers

To manage authorized peers:

1. Select **Configuration > ADN > Manager > Peer Authorization**.



2. To manually add peers that are authorized to join the ADN:
 - a. Click **Add**. The Add ADN Peers dialog is displays.
 - b. Enter the device IDs for the ADN nodes you want to authorize and then click **OK**. To find the device ID for a node, see the **Extracted Device ID** field on that node (on the **ADN > General > Device Security** tab).
3. To remove a peer that was previously authorized to join the ADN, select the node from the **Approved Peers** list and then click **Remove**. If a peer is deleted from the approved list, the ADN manager broadcasts a `REJECT-PEER` to all peers to delete this peer and terminate any existing ADN connections to it. No new connections are routed through the deleted ADN peer.

Note: If you remove a peer and then want it to rejoin the ADN, you must reconnect the peer to the ADN manager(s). Select **Configuration > ADN > General > Reconnect to Managers**.


4. Click **Apply**.

Approving a Peer

To approve a peer:

If a peer is configured to contact the ADN manager on startup but has not been added to the approved list, the ADN manager adds the peer to the list of pending peers if the **Allow Pending Peers** option is selected. You must manually move a peer from the Pending Peers list to the Approved Peers list on both the Primary ADN Manager and the Backup ADN Manager as follows:

1. Select **Configuration > ADN > Manager > Pending Peers**.



2. Select the **Allow Pending Peers** option.
3. To manage pending peers:
 - Highlight a peer and click **Accept** or **Reject**; alternatively, you can select or reject all peers in the list by clicking **Accept All** or **Reject All**. If accepted, the peer moves to the **Approved** list; if not, it is dropped from the **Pending Peers** list.
 - You can also leave peers in the pending list by not selecting them or selecting them and clicking **Mark Pending**.
4. Click **Apply**.

Section D: Configuring Load Balancing

This section discusses the following topics:

- ❑ ["Introduction to Load Balancing"](#)
- ❑ ["Configuring Transparent Load Balancing" on page 746](#)
- ❑ ["Configuring Explicit Load Balancing" on page 747](#)

Introduction to Load Balancing

The way you configure load balancing depends on whether you are using explicit or transparent tunnels as described in the following sections:

- ❑ ["Configuring Transparent Load Balancing" on page 746](#)
- ❑ ["Configuring Explicit Load Balancing" on page 747](#)

Configuring Transparent Load Balancing

There are two ways to configure transparent load balancing as described in the following sections:

- ❑ ["Using a ProxySG as a Transparent Load Balancer" on page 746](#)
- ❑ ["Using a WCCP Router or L4 Switch as a Load Balancer" on page 747](#)

Using a ProxySG as a Transparent Load Balancer

When you configure transparent load balancing using a ProxySG as the load balancer, the ProxySG that is designated as the load balancer is deployed in-path and therefore receives all traffic destined for WAN optimization. This ProxySG then determines the ProxySG to which to send each packet for optimization. You can optionally designate the ProxySG as a dedicated load balancer, meaning that it does not participate in ADN tunnel connections.

The ProxySG can intercept IPv4 or IPv6 connections and load balance these connections in a connection forwarding cluster. Note that all ProxySG appliances in the forwarding cluster must be able to handle the address type (IPv4 vs. IPv6) of the connection.

To configure a ProxySG as a transparent load balancer:

1. Deploy the load-balancing ProxySG in-path so that it can transparently intercept all traffic.
2. Enable load balancing on all peers by selecting **Configuration > ADN > Tunneling > Load Balancing**, and selecting the **Enable Load Balancing** option.
3. (Optional) If you do not want this ProxySG to participate in any ADN tunnels (that is, you want it to act as a dedicated load balancer), select **Act as load balancer only**. This ProxySG is still part of the ADN and must still connect to the ADN manager(s).

4. Put all ADN peers into a forwarding connection cluster. For more information, see ["TCP Connection Forwarding"](#) on page 857.
5. (Optional) Set the same group name on all of the peers in the cluster.

Using a WCCP Router or L4 Switch as a Load Balancer

When you configure transparent load balancing using a WCCP router or L4 switch as the load balancer in an IPv4-only network, the WCCP router or switch redirects traffic to a ProxySG in the load balancing group. The ProxySG that receives the redirected traffic from the router or switch then determines which ProxySG in the group should handle the traffic.

The procedure to configure a WCCP router or L4 switch as a load balancer is similar to the procedure for using a ProxySG as the load balancer, except that you must also define the WCCP router or L4 switch configuration on each node in the cluster.

Note: Symantec does not currently support WCCP on an IPv6 network.

To configure transparent load balancing using a WCCP router or L4 switch:

1. Enable load balancing on all peers by going to **Configuration > ADN > Tunneling > Load Balancing**, and selecting the **Enable Load Balancing** option.
2. Put all ADN peers into a connection forwarding cluster. For more information, see ["TCP Connection Forwarding"](#) on page 857.
3. (Optional) Configure each box in the cluster with the same load-balancing group name.
4. Configure WCCP on each peer and on the WCCP router. For detailed information on configuring WCCP, refer to the *WCCP Reference Guide*.

Configuring Explicit Load Balancing

There are two ways to configure explicit load balancing as described in the following procedures:

- ❑ ["Configuring Explicit Load Balancing Using Server Subnets"](#) on page 747
- ❑ ["Configuring Explicit Load Balancing Using an External Load Balancer"](#) on page 748

Configuring Explicit Load Balancing Using Server Subnets

When using the server subnet method to achieve explicit load balancing, you simply place multiple ProxySG appliances in front of the same IPv4 and/or IPv6 server subnet. You then configure the server subnet on each ADN peer in the group. If multiple Concentrator peers are configured as Internet gateways, Branch peers will choose only those Concentrator peers that contain at least one address of the same family as the destination address.

To configure explicit load balancing using server subnets:

1. On each peer in the group, select **Configuration > ADN > Routing**.
2. Click **Add**.
3. Add the IPv4 or IPv6 subnet route to be advertised by the ADN manager and then click **OK**.

For detailed information about configuring server subnets, see "[Advertising Server Subnets](#)" on page 729.

Configuring Explicit Load Balancing Using an External Load Balancer

Using an external load balancer provides more control than using server subnets alone for external load balancing. However, it requires more configuration on each node. Only use a virtual IP (VIP) address type (IPv4 vs. IPv6) that can be reached from all Branch peers.

To configure explicit load balancing using an external load balancer:

1. On each ADN peer, define the subnets to be advertised on the load balanced subnets. See "[Advertising Server Subnets](#)" on page 729.
2. On each ADN node, configure the VIP of the external load balancer by selecting **Configuration > ADN > Tunneling > Load Balancing** and entering the IPv4 or IPv6 address in the **External VIP** field.
 - In a homogeneous ADN in which Branch peers can reach only an IPv4 VIP, configure an IPv4 VIP on the Concentrator peers.
 - In a homogeneous ADN in which Branch peers can reach only an IPv6 VIP, configure an IPv6 VIP on the Concentrator peers.
 - In a heterogeneous ADN in which some Branch peers support only one type of address, configure a VIP of the type that is supported by all Branch peers in the ADN. For example, if the ADN contains one or more Branch peers that are only IPv4 capable, then the other Branch peers that are IPv6 capable should still be configured with an IPv4 address to reach an IPv4 VIP.
 - In an ADN where all Branch peers are capable of connecting to an IPv4 or IPv6 VIP, you can choose to configure the VIP as either version.
3. Click **Apply**.

Section E: Configuring Advanced ADN Settings

The following sections describe optional ADN configuration tasks. These tasks are not required for basic ADN setup, but you may choose to configure these options in some situations.

- ❑ "Configuring Adaptive Compression" on page 749
- ❑ "Configuring an ADN Node as an Internet Gateway" on page 750
- ❑ "Modifying the TCP Window Size" on page 751
- ❑ "Configuring the Byte-Cache Dictionary Size" on page 752
- ❑ "Deleting ADN Peers" on page 756


Configuring Adaptive Compression

Adaptive compression enables the ProxySG to adjust its compression level based on CPU usage. When adaptive compression is enabled, the ProxySG will automatically increase its compression level when CPU usage is low and decrease its compression level when CPU usage is high.

All ProxySG platforms that are manufactured or remanufactured with a SGOS 6.2 or later release have adaptive compression enabled by default. In the case of an upgrade to SGOS 6.2 or later, the setting matches the configuration before the upgrade. For example, if adaptive compression was disabled in SGOS 6.1, it will be disabled after upgrading to SGOS 6.2 or later.

To enable adaptive compression:

1. Select **Configuration > ADN > Byte Caching**.



2. Select (or deselect) the **Enable adaptive compression** option to enable (or disable) adaptive compression
3. Click **Apply**.

Configuring an ADN Node as an Internet Gateway

You can configure an ADN node as an Internet gateway for IPv4 or IPv6 addresses. Subnets that should not be routed to the Internet gateway can be configured as exempt subnets.


In explicit deployments:

- ❑ An IPv6-only Concentrator peer will not be advertised as the Internet gateway for a node that is running an older (pre-6.2.4) version of software.
- ❑ An IPv4-only Branch peer running SGOS 6.2.4 or higher will not use an IPv6-only Concentrator peer as an Internet gateway.
- ❑ Similarly, an IPv6-only Branch peer will not use an IPv4-only Concentrator peer as an Internet gateway.

Note: You can also configure the exempt subnet capability through policy that allows you to disable ADN tunnels for specific connections. For more information, refer to *Content Policy Language Guide*.

To enable this peer as an Internet gateway:

1. Select **Configuration > ADN > Routing > Internet Gateway**.



2. Select **Enable this SG as an Internet Gateway for all subnets except the following**.
3. Click **Add**. The Add IP/Subnet dialog displays.
4. Define each subnet to be exempted, and then click OK:

Note: Some subnets are on the exempt list by default (for example, 10.0.0.0/8 and fe80::/10). Verify these default exempt defaults do not affect the configuration in your environment.

- **IP address:** Enter an IPv4 or IPv6 address.
 - **Prefix length or subnet mask:** Specify the prefix length (for IPv6) or subnet mask (for IPv4).
5. Repeat steps 3 and 4 for each subnet.
 6. Click **Apply**.

Modifying the TCP Window Size

TCP window size is the number of bytes that can be buffered on a system before the sending host must wait for an acknowledgement from the receiving host. The TCP window size for ADN tunnel connections is set and updated automatically, based on current network conditions and on the receiving host's acknowledgement. In most situations, you do not need to modify the TCP window size. You might need to modify it only if your network environment has intervening network equipment that makes the delay appear lower than it actually is. These environments are sometimes found on satellite links that have high bandwidth and high delay requirements. In this case, the automatically adjusted window size would be smaller than optimal.

Note: If you know the bandwidth and round-trip delay, you can compute the value to use as, roughly, $2 * \text{bandwidth} * \text{delay}$. For example, if the bandwidth of the link is 8 Mbits/sec and the round-trip delay is 0.75 seconds:

$$\text{window} = 2 * 8 \text{ Mbits/sec} * 0.75 \text{ sec} = 12 \text{ Mbits} = 1.5 \text{ Mbytes}$$

The setting in this example would be 1500000 bytes. This number goes up as either bandwidth or delay increases, and goes down as they decrease.

You can decrease or increase the window size based on the calculation; however, decreasing the window size below 64Kb is not recommended.

The window-size setting is a maximum value; the normal TCP/IP behaviors adjust the window-size setting downward as necessary. Setting the window size to a lower value might result in an artificially low throughput.

To modify the TCP window size:

1. Select **Configuration > ADN > Tunneling > Network**.
2. In the **TCP Settings** section of the window select **Manual override** and then enter the window size in the text box. The configurable range is between 8 Kb and 4 MB (8192 to 4194304), depending on your bandwidth and the round-trip delay. Setting sizes below 64 Kb are not recommended.
3. Click **Apply**.

Configuring the Byte-Cache Dictionary Size

When byte caching is in effect for an application, byte sequences in traffic flows are replaced with reference tokens. The byte sequences are stored in a *byte-cache dictionary* on a pair of ProxySG appliances at each end of the WAN. When a matching byte sequence is requested again, the ProxySG transmits a token instead of the byte sequence.

If a ProxySG forms tunnel connections with multiple ProxySG appliances, it will have a separate byte-cache dictionary for each peer. Because these dictionaries will need to share the available disk space, the ProxySG automatically determines how much disk space to allocate to each peer based on the traffic history of each peer and the effectiveness of byte caching on the applications that are being accelerated on that peer. The peers are then ranked and disk space is allocated based on these rankings.

Note: Peers that are using an SGOS version prior to 5.3 do not support persistent byte-cache, so GZIP-only mode is used on these nodes. Therefore, they are not ranked unless you have manually sized their dictionaries.

In some instances you may want to manually set the size of a peer dictionary. For example, suppose you have a mission critical application that you want to accelerate using byte caching. If byte caching isn't as efficient for this application as for other applications accelerated by other peers, the peer may not be allocated any dictionary space or may be allocated a small dictionary. If you want to ensure that this mission-critical application can use byte caching, you might want to manually resize its dictionary. Keep in mind that any manually-sized peers are ranked above all other peers. In addition, the automatic dictionary sizing feature is no longer in effect for this peer, so you should not use this feature unless absolutely necessary.

Note: You cannot reduce the space available for byte caching to below the total size of all manually sized dictionaries. You also cannot assign a size to a dictionary that would cause the total size of all manually sized dictionaries to exceed the space available for byte-caching.

Because a byte-cache dictionary is shared between two peers, any time you make a change to the dictionary on one peer, you must make the same change on the other peer. For example, if you manually size a dictionary to a particular size on one peer, you must change the other peer's dictionary to manual and set it to the same size. There are two ways to manually resize the byte-cache dictionaries depending on whether or not the peer already has a dictionary established:

- ❑ If a dictionary already exists for the peer, see "[Manually Resizing the Byte Cache Dictionaries From the Statistics Tab](#)" on page 753.
- ❑ If the peer does not yet have an established dictionary, see "[Manually Resizing Byte Cache Dictionaries from the Byte Caching Tab](#)" on page 754.

Manually Resizing the Byte Cache Dictionaries From the Statistics Tab

This section discusses how to manually resize byte cache dictionaries from the Statistics tab. To manually resize dictionaries from the Byte Caching tab instead, see ["Manually Resizing Byte Cache Dictionaries from the Byte Caching Tab"](#) on page 754.

For more information about these options, see ["Configuring the Byte-Cache Dictionary Size"](#) on page 752.

To manually resize byte cache dictionaries from the Statistics tab:

1. Select **Statistics > ADN History > Peer Dictionary Sizing**.


Rank ▲	Peer ID	Peer IP	Byte Cache Score	Peer Traffic (GB/Day)	Fill Rate (GB/Day)	Recommended Dict Size (GB)
1	106080041	10.9.80.136	14354897253	49.2820	0.0000	4.3022

The **Peer Dictionary Sizing** tab displays the following statistics for each peer.

- **Rank:** The ranking of a peer's dictionary. Manually-configured peers have a higher rank than dynamically-configured peers.
- **Peer ID:** The serial number of the device.
- **Peer IP:** The IPv4 or IPv6 address of the device, if it is connected.
- **Byte Cache Score:** The score of this peer relative to other peers. Score is calculated based on the traffic history and byte-caching efficiency of the peer.
- **Peer Traffic (GB/Day):** The average amount of pre-byte-cache traffic per day.
- **Fill Rate (GB/Day):** The average amount of data put into the dictionary per day over the last week.
- **Recommended Dict Size (GB):** The dictionary size the Blue Coat appliance recommends, based on the peer traffic over the last week.
- **Actual Dict Size (GB):** The actual size of the dictionary.

Note: You can also delete a peer from this tab. For more information, see ["Deleting ADN Peers"](#) on page 756.

2. Select the peer for which you want to resize the dictionary and click **Edit**. The **Edit Peer** dialog displays.



3. To set the dictionary size for the selected peer, select the **Manual Re-size** radio button and enter the desired dictionary size value (in megabytes).
4. Click **OK**. The peer dictionary is resized immediately. You must manually size the corresponding peer's dictionary to the same size.


Manually Resizing Byte Cache Dictionaries from the Byte Caching Tab

This section discusses how to manually resize byte cache dictionaries from the Byte Caching tab. To manually resize dictionaries from the Statistics tab instead, see ["Manually Resizing the Byte Cache Dictionaries From the Statistics Tab"](#) on page 753.

For more information about these options, see ["Configuring the Byte-Cache Dictionary Size"](#) on page 752.

To manually size byte cache dictionaries from the Configuration > ADN > Byte Caching tab:

1. Select **Configuration > ADN > Byte Caching**.




Note: You can also enable or disable adaptive compression from this tab. For more information, see ["Configuring Adaptive Compression"](#) on page 749.

- To change the total disk space available for all byte-cache dictionaries, change the percentage in the **Maximum disk space to use for byte caching** field.

The **Max disk usage range should be between 5 and 80 percent of x GB** indicates how much of the existing disk space can be used for byte caching.


- Click **New**. The **Create Manual Dictionary Sizing** dialog box displays.



- Enter the peer ID (serial number) of the device for which you want to manually size the dictionary.

5. Enter the new value in megabytes in the **Size** field or select the **Disable Byte Caching** radio button to disable byte caching for this peer.
6. Click **OK**.
7. Click **Apply**. The peer is added to the manually configured dictionary sizing list and is ranked among the other manually sized peers at the top of the dictionary byte cache table. You must manually size the corresponding peer's dictionary to the same size.

Note: If you enter an invalid value, an error message displays when you click **Apply**. The error message displays the maximum disk space you can allocate to the manually-sized dictionary.



To change a manually sized dictionary to an automatically sized dictionary:

1. Select **Configuration > ADN > Byte Caching**. This tab displays all peers that have manually sized dictionaries.
2. Select the peer you want to convert from manual dictionary sizing to automatic dictionary sizing and click **Delete**.
3. Make the same changes on the corresponding peer. For example, if you changed this peer's byte-cache dictionary from manually-sized to automatically-sized, you must also change the corresponding peer's dictionary to automatically-sized.

Deleting ADN Peers

The ProxySG allocates space in its byte-cache dictionary for each ADN peer that forms a tunnel connection with it. If the maximum number of ADN peers is reached (the maximum number of peers that is supported depends on the size of the system), any new peer that forms a tunnel connection with the ProxySG cannot be allocated dictionary space. Therefore, traffic to and from this peer cannot be accelerated using byte caching; instead only GZIP compression is used.

To prevent this, each day after it updates its traffic history the ProxySG automatically deletes peers that meet the following criteria:

- The dictionary for the peer is empty and is automatically sized
- The peer has been idle for at least eight days
- There is no active connection (data or control) with the peer

Note: Automatic peer deletion occurs at 3:05 AM local standard time. If you change the time zone you must reboot the appliance in order for ADN to use the new time.

As long as your system is sized properly, the automatic peer deletion process will prevent you from reaching the maximum number of peers. However, there may be times when you want to manually delete a peer that you know is no longer valid (and is therefore taking up dictionary space unnecessarily) and that will not get deleted automatically, either because its dictionary is manually sized or because it has not yet been idle for at least 8 days .

Keep in mind that even if you delete a peer, it can be accepted as a peer again if it forms a tunnel connection later.

Note: You cannot delete ProxyClient ADN peers from the Management Console; you must use the CLI instead.

To manually delete an ADN peer:

1. Select **Statistics > ADN History > Peer Dictionary Sizing**.
2. Select the peer you want to delete and click **Delete**. All ProxyClient peers are displayed in a single line and cannot be deleted. You must delete ProxyClient peers using the CLI.
3. When prompted to confirm the deletion, click **Yes**.

Note: Sometimes, the system may be unable to delete a peer if it is performing internal maintenance tasks. If this happens, try deleting the peer again later.


Section F: Monitoring the ADN

After you have configured and enabled ADN, you can review various ADN history and statistics as follows:

- ❑ "Reviewing ADN History" on page 758
- ❑ "Reviewing ADN Active Sessions" on page 759
- ❑ "Monitoring Adaptive Compression" on page 760
- ❑ "Reviewing ADN Health Metrics" on page 762

Reviewing ADN History

Review the ADN history by selecting **Statistics > ADN History**.



You can view either usage statistics or gain statistics (by clicking the **Gain** tab) and either **Unoptimized Bytes** or **Optimized Bytes** through the pie charts on the right side.

The left side of the tab represents optimized and unoptimized bytes trend graphs for the selected peer or all peers; hovering the cursor over the graph displays statistics in numeric form. For definitions of each of the statistics in the tool tips, see "Viewing Bandwidth Details for Proxies or Services" on page 668.

The right-side pie chart represents optimized and unoptimized bytes for all peers. The rows in the table below the graphs represent ADN peers and columns representing various aspects of the ADN peers:

Note: All ProxyClient peers are combined and shown on one row. For more information on ProxyClient refer to the *ProxyClient Configuration and Deployment Guide*.

- ❑ **Peer ID:** ID of the peer.
- ❑ **Peer IP:** IPv4 or IPv6 address of the peer.
- ❑ **Optimized Bytes:** Data that has been byte-cached and/or compressed.
- ❑ **Unoptimized Bytes:** Data that is to be byte-cached or compressed and data that has been *un*-byte-cached or decompressed.
- ❑ **Savings:** The percentage of data that did NOT have to be sent over the WAN because of object and byte caching, protocol optimization, and compression. Moving the cursor over the **Savings** column value displays tool-tip information.

Selecting any row in the table changes the trend graph at top left and display graphs for the selected peer. If you select the last row, which displays totals, the trend graph at top left reflects the cumulative data. Changing the duration (using the **Duration** drop-down list) changes the graph accordingly.

Reviewing ADN Active Sessions

You can view active ADN inbound connections through the **Statistics > Sessions > Active Sessions > ADN Inbound Connections**. Information from the **ADN Inbound Connections** tab can be used for diagnostic purposes.

These connections are not persistent. When a connection completes, the statistics for that connection no longer display.

Client	Server	Peer	Duration	Unopt. Bytes	Opt. Bytes	Savings	C	BC	E	Tunnel Type	Detail
10.9.45.130	10.9.45.129:80	10.9.45.128	46 sec	121,816,093	445,768	99.63%				Explicit	
2001:418:9804:111::169	2001:418:9804:111::161:445	2001:418:9804:111::166	3.3 min	177,439	11,403	93.57%				Explicit	

Terminate Connection Download

Total displayed connections: 2

You can filter on a number of variables, including client, server, or peer IP address; server port, or none (shown above). You can also limit the number of connections being displayed to the *n* most recent.

Note: You must press **Show** each time you change display options or if you want to refresh the page.

You can terminate an active ADN inbound connection or you can download session details.

- ❑ To terminate an ADN inbound connection, select the session in the list and click **Terminate Connection**.
- ❑ To download details about all connections as a text file that you can open in a spreadsheet program, click **Download**. All of the connections in the list are downloaded.

Each connection has the following details.

Client: The IP address of the system that is being sent through the ProxySG over ADN connections.

Server: The IP address of the server to which you are connecting: CNN, for example, or Google.

Peer: The downstream ProxySG or ProxyClient. The type of address (IPv4 vs. IPv6) indicates the type of tunnel. For example, if the peer address is 2001:418:9804:111::169, it is an IPv6 tunnel. Or if the peer address is 10.9.45.129, it is an IPv4 tunnel.

Duration: The length of time the connection has been active.

Unopt. Bytes: The amount of data served to/from the server prior to or subsequent to ADN optimization.

Opt. Bytes: The amount of compressed/byte-cached data sent to/received from the downstream ProxySG/ProxyClient.

Savings: A relative percentage of bandwidth savings on the WAN link.

Compression: Whether gzip compression is active in either direction on that tunnel.

Byte Caching: Whether byte caching is active in either direction on that tunnel.

Encryption: Whether encryption is active in either direction on that tunnel.

Tunnel Type: One of the following: Explicit, Transparent, or Client.


Monitoring Adaptive Compression

When adaptive compression is enabled, the ProxySG determines whether to increase or decrease the compression level based on CPU usage. When extra CPU is available, it will adapt compression to use these additional resources, resulting

in higher CPU usage. Therefore, when this feature is enabled, you should monitor adaptive compression in addition to CPU usage statistics when making capacity planning decisions.

To monitor adaptive compression:


1. Select **Statistics > ADN History > Adaptive Compression**. A graph detailing adaptive compression over the last hour is displayed. The bars on the graph display in three colors, indicating if or how compression has been adapted:
 - **Green**—Indicates that the ProxySG has adapted compression to operate at a higher level to take advantage of available CPU resources.
 - **Yellow**—Indicates that compression is operating at the ideal level.
 - **Red**—Indicates that the ProxySG has adapted compression to operate at a lower level due to a lack of CPU resources; any additional load may impact performance. If you notice that adaptive compression displays red consistently, your appliance may be undersized; consider a hardware upgrade.



2. (optional) To monitor adaptive compression over a different time range, select a new time range from the **Duration** drop-down list.

Reviewing ADN Health Metrics

To view ADN health metrics, select **Statistics > Health Monitoring> Status**.



The screenshot shows a web interface with three tabs: General, Licensing, and Status. The Status tab is selected and displays a table of ADN health metrics. The table has three columns: Metric, Value, and State. The metrics include disk status, motherboard and CPU temperatures, fan speeds, bus voltages, and ADN connection status. The 'ADN Connection Status' row is highlighted in blue.

Metric	Value	State
Disk 1 Status	present	OK
Disk 2 Status	present	OK
Disk 3 Status	present	OK
Disk 4 Status	present	OK
Motherboard temperature	30.8 degrees C	OK
CPU 1 temperature	40.0 degrees C	OK
Fan 1 speed	7314.18 RPM	OK
Fan 2 speed	7602.44 RPM	OK
Fan 3 speed	7483.04 RPM	OK
+3.3V bus voltage	3.30 Volts	OK
+5V bus voltage	4.95 Volts	OK
+5V standby voltage	5.05 Volts	OK
+12V bus voltage	12.04 Volts	OK
+3.3V bus voltage 2 (Vcc)	3.30 Volts	OK
CPU core voltage	1.33 Volts	OK
ADN Connection Status	Connected to closed ADN network	OK
ADN Manager Status	Not a manager	OK
Health Check Status	OK	OK

The **Status** tab displays ADN health statistics for the following metrics:

- ❑ ADN Connection Status
- ❑ ADN Manager Status

The following table describes the possible values for each metric, which you can use for diagnostic and debugging purposes.

Table 32–2 ADN Health Metrics

Metric	Value	Description	State
ADN Connection Status	Connected to closed ADN network	The ADN peer is connected to the ADN manager, ready to receive any route/peer updates. If a backup manager exists, this state indicates the peer is connected to both managers.	OK
	Connected to open ADN network	ADN is enabled on the peer and is ready to form connections with other peers	OK
	Functionality Disabled	ADN functionality is not enabled.	OK
	Not operational	ADN functionality is not operational yet — components are starting up or shutting down.	OK
	Open ADN	The node is operating in Open ADN mode	OK
	Connection Approved	The ADN peer has been approved to connect to the ADN manager.	OK
	Connecting	The ADN peer is in process of connecting to ADN manager.	OK
	Partially connected to closed ADN network	The ADN peer is connected to one ADN manager but not the other.	Warning
	Partially connected to open ADN network	The ADN peer is connected to one ADN manager but not the other.	Warning
	Mismatching Approval Status	The ADN peer is approved by the current active ADN manager but is rejected by the backup manager. This warning only exists if a backup ADN manager is configured.	Warning
	Approval Pending	The ADN peer is awaiting a decision from the active ADN manager for the peer's request to join the ADN.	Warning
	Disconnected	The ADN peer is not connected to the ADN manager and cannot receive route/peer information. If a backup manager is configured, this state indicates the peer is disconnected from both manager peers.	Critical
	Connection Denied	The ADN peer is rejected by the ADN managers in the peer's request to join the ADN.	Critical

Table 32–2 ADN Health Metrics

Metric	Value	Description	State
ADN Manager Status	Not an ADN manager	The ADN peer is not an ADN manager.	OK
	No Approvals Pending	All ADN peers that are requesting to join the network are already on the approved list.	OK
	Approvals Pending	ADN peers are requesting to join the network. The approvals are made by the administrator.	Warning

Section G: Related CLI Syntax to Configure an ADN

- ❑ To enter configuration mode:

```
SGOS#(config) adn
SGOS#(config adn)
```

Note: For detailed information on these commands, refer to the *Command Line Interface Reference*.

- ❑ The following subcommands are available:

```
SGOS#(config adn) {enable | disable}
SGOS#(config adn) exit
SGOS#(config adn) byte-cache
    SGOS#(config adn byte-cache) adaptive-compression {enable | disable}
    SGOS#(config adn byte-cache) delete-peer peer-id [force]
    SGOS#(config adn byte-cache) max-disk-usage percentage
    SGOS#(config adn byte-cache) peer-size peer-id {size_in_megabytes | auto | none}
    SGOS#(config adn byte-cache) exit
    SGOS#(config adn byte-cache) view
SGOS#(config adn) load-balancing
    SGOS#(config adn load-balancing) {enable | disable}
    SGOS#(config adn load-balancing) exit
    SGOS#(config adn load-balancing) external-vip IP_address
    SGOS#(config adn load-balancing) group group_name
    SGOS#(config adn load-balancing) load-balance-only {enable | disable}
    SGOS#(config adn load-balancing) no {external-vip | group}
    SGOS#(config adn load-balancing) view
SGOS#(config adn) manager
    SGOS#(config adn manager) backup-manager {IP_address [ID] | self | none}
    SGOS#(config adn manager) exit
    SGOS#(config adn manager) open-adn {enable | disable}
    SGOS#(config adn manager) port port_number
    SGOS#(config adn manager) primary-manager {IP_address [ID] | self | none}
    SGOS#(config adn manager) secure-port secure_port_number
    SGOS#(config adn manager) view [approved-peers | backup-manager-id | pending-peers | primary-manager-id]
SGOS#(config adn manager) approved-peers
    SGOS#(config adn approved-peers) add peer-device-ID
    SGOS#(config adn approved-peers) exit
    SGOS#(config adn approved-peers) remove peer-device-ID
    SGOS#(config adn approved-peers) view
```

```
SGOS#(config adn manager) pending-peers
    SGOS#(config adn pending-peers) {accept | reject}
    SGOS#(config adn pending-peers) {enable | disable}
    SGOS#(config adn pending-peers) exit
    SGOS#(config adn pending-peers) view
SGOS#(config adn) routing
    SGOS#(config adn routing) exit
    SGOS#(config adn routing) prefer-transparent {enable | disable}
    SGOS#(config adn routing) view
    SGOS#(config adn routing) advertise-internet-gateway
        SGOS#(config adn routing advertise-internet-gateway) {disable | enable}
        SGOS#(config adn routing advertise-internet-gateway) exempt-subnet {add {subnet_prefix[/prefix_length]} clear-all | remove {subnet_prefix[/prefix_length]} | view}
        SGOS#(config adn routing advertise-internet-gateway) exit
        SGOS#(config adn routing advertise-internet-gateway) view
    SGOS#(config adn routing) server-subnets
        SGOS#(config adn routing server-subnets) add subnet_prefix [/ prefix length]
        SGOS#(config adn routing server-subnets) clear-all
        SGOS#(config adn routing server-subnets) remove subnet_prefix [/ prefix length]
        SGOS#(config adn routing server-subnets) exit
        SGOS#(config adn routing server-subnets) view
SGOS#(config adn) security
    SGOS#(config adn security) authorization {enable | disable}
    SGOS#(config adn security) exit
    SGOS#(config adn security) manager-listening-mode {plain-only | plain-read-only | secure-only | both}
    SGOS#(config adn security) no ssl-device-profile
    SGOS#(config adn security) secure-outbound {none | secure-proxies | all}
    SGOS#(config adn security) ssl-device-profile profile_name
    SGOS#(config adn security) tunnel-listening-mode {plain-only | secure-only | both}
    SGOS#(config adn security) view
SGOS#(config adn) tunnel
    SGOS#(config adn tunnel) connect-transparent {enable | disable}
    SGOS#(config adn tunnel) exit
    SGOS#(config adn tunnel) preserve-dest-port {enable | disable}
    SGOS#(config adn tunnel) port port_number
    SGOS#(config adn tunnel) reflect-client-ip (deny | allow | use-local-ip)
    SGOS#(config adn tunnel) secure-port secure_port_number
    SGOS#(config adn tunnel) tcp-window-size {auto | window_size_in_bytes}
    SGOS#(config adn tunnel) view
```

Section H: Policy

The following gestures can be used for WAN optimization from either the VPM or CPL.

Note: For more information on using the VPM or CPL to configure policy, refer to *Visual Policy Manager Reference* or *Content Policy Language Guide*.

- ❑ `adn.server(yes | no)` (This property overrides all other routing and intercept decisions made by ADN based on configuration and routing information.)
- ❑ `adn.server.optimize(yes | no)`
- ❑ `adn.server.optimize.inbound(yes | no)`
- ❑ `adn.server.optimize.outbound(yes | no)`
- ❑ `adn.server.optimize.byte-cache(yes | no)`
- ❑ `adn.server.optimize.inbound.byte-cache(yes | no)`
- ❑ `adn.server.optimize.outbound.byte-cache(yes | no)`
- ❑ `adn.server.optimize.compress(yes | no)`
- ❑ `adn.server.optimize.inbound.compress(yes | no)`
- ❑ `adn.server.optimize.outbound.compress(yes | no)`
- ❑ `adn.server.dscp`

Section I: Troubleshooting

You can troubleshoot your ADN several ways:

- ❑ through the `test adn` diagnostics command
- ❑ through viewing the ADN configuration

Each of these tools can provide information about the ADN and suggest reasons for the network failure.

Using the Test ADN Diagnostics Command

the `test adn` command is used to test connectivity from one ProxySG to an IPv4 or IPv6 server on a specified port. This test also can be done with an ADN port to test the success or failure of a ProxySG connection to an ADN peer.

The command provides details of its success or failure.

Transparent ADN: Success

```
Blue Coat SG200 Series# test adn 192.168.0.222 80
connecting to 192.168.0.222:80...succeeded!
Diagnostics
Route decision   : Connect Transparently
Route reason     : ADN transparent due to no explicit route
Route policy     :
Connect result   : Success
Remote peer      : 207060009
Local Addr       : 192.168.0.121:64881
Peer Addr        : 192.168.0.222:80
```

Notes

- ❑ If the Branch ADN peer is able to successfully reach the OCS by forming a transparent ADN tunnel, you will see the Success messages shown above.
- ❑ The **Remote Peer** is the device ID (serial number, in this case) of the remote ProxySG the `test adn` command found. When last peer detection is enabled on intermediate concentrators and you issue the `test adn` command from the Branch peer, the **Remote Peer** should be the last qualified peer, such as the ProxySG closest to the OCS.
- ❑ The **Local Addr** is the originating system.
- ❑ The **Peer Addr** shows either the server IP address (for transparent tunnels, as in this example) or the ProxySG IP address (for explicit or translucent tunnels).

Transparent ADN: Success but no Upstream ADN Connection

```
Blue Coat SG200 Series# test adn 192.168.0.222 80
Connecting to 192.168.0.222:80...succeeded!
Diagnostics
Route decision   : Attempted Transparent but went Direct
Route reason    : ADN transparent due to no explicit route
Route policy    :
Connect result   : Success
Peer Addr       : 192.168.0.222:80
```

Notes

- ❑ Because no ADN connection existed, the **Route decision** indicates what happened:
 - The `test adn` command went directly to the server.
 - **Success** in this case refers to the successful connection to the server but not through an ADN connection.
 - Remote peer device ID and local address information were not available.

Explicit ADN: Success

```
Blue Coat SG200 Series# test adn 192.168.0.222 80
Connecting to 192.168.0.222:80...succeeded!
Diagnostics
Route decision   : Connect Explicitly
Route reason    : ADN explicit route found
Route policy    :
Explicit routes found:
    Peer (207060009) ip#0: 192.168.0.122, ports: 3035,3037 Connect
result   : Success
Remote peer   : 207060009
Local Addr    : 192.168.0.121:53892
Peer Addr     : 192.168.0.122:3035
```

Notes

- ❑ The **Remote Peer** is the device ID (serial number, in this case) of the remote ProxySG the `test adn` command found.
- ❑ The **Local Addr** is the originating system.
- ❑ The **Peer Addr** is the IP address of the remote peer (for explicit or translucent tunnels) or the IP address of the server (for transparent tunnels).

Explicit ADN: The Upstream Device is not Functioning

```
Blue Coat SG200 Series# test adn 192.168.0.222 80
Connecting to 192.168.0.222:80...failed with error : 5!
Diagnostics
Route decision : Connect Explicitly
Route reason   : ADN explicit route found
Route policy   :
Explicit routes found:
    Peer (207060009) ip#0: 192.168.0.122, ports: 3035,3037
Connect result : Failure
Failure reason  : Socket internal error
Network error   : Socket error(5)
Local Addr     : 192.168.0.121:53892
Peer Addr      : 192.168.0.122:3035
```

Notes

- ❑ For an explicit connection, the local IP address is displayed even if a connection cannot be established.

Error Codes

Table 32–3 Error Codes

Error Code	Description
5	Networking Input/output error
50	Network is down
51	Network is unreachable
52	Network dropped connection on reset
53	Software caused connection abort
54	Connection reset by peer
55	No buffer space available
56	Socket is already connected
57	Socket is not connected
58	Can't send after socket shutdown
59	Too many references: can't splice
60	Operation timed out
61	Connection refused

Showing the ADN Configuration

You can view the entire ADN configuration through the **show adn** CLI command. Also, you can use the **show adn** subcommands to view specific parts of the ADN configuration. This section describes the **show adn** subcommands.

- ❑ **ADN Manager Configuration:** The manager configuration shows the primary and backup managers, ports, and where approved devices connect from.

```
SGOS# show adn manager
Primary manager:          self
Backup manager:           10.9.59.243 2505060056
Port:                     3035
Secure port:              3037
Approved device           Connecting from
2505060056 10.25.36.48
Allow pending devices:    enabled
Pending device            Connecting from
```

- ❑ **Tunnel Configuration:** The tunnel configuration displays connection information for this device.

```
SGOS# show adn tunnel
Port:                     3035
Secure port:              3037
proxy-processing http:    disabled
connect-transparent:      enabled
preserve-dest-port:       disabled
TCP window size:          auto
reflect-client-ip:        use-local-ip
```

- ❑ **Load Balance Configuration:** The load balance configuration displays the Load Balance information for this device.

```
SGOS# show adn load-balancing
Load Balancing Configuration:
Load-balancing:            disabled
Load-balancing Group:      <none>
Load-balance only mode:    disabled; will take traffic
External VIP:              none
```

- ❑ **Routing Table:** The routing table section shows the advertised subnets for this device. The routing table is only populated if explicit ADN is used.

```
SGOS# show adn routing
Prefer Transparent:        disabled
Internet Gateway:          enabled
Exempt Server subnet:      10.0.0.0/8
Exempt Server subnet:      172.16.0.0/12
Exempt Server subnet:      192.168.0.0/16
Server subnet:             10.25.36.0/24
```

- ❑ **Security Configuration:** This section displays security information about the device.

```
SGOS# show adn security
Ssl-device-profile:      bluecoat-appliance-certificate (Device-id:
4605060001)
Manager-listening mode:  both
Tunnel-listening mode:  both
Authorization:           enabled
Secure-outbound:         secure-proxies
```

- ❑ **Byte Cache Configuration:** This section shows the percentage of disk space you are allowing this peer to use for byte caching. The recommended range is also displayed. For more information on the byte-caching CLI tables that are displayed as part of the byte-cache configuration output, continue with the next section.

```
SGOS# show adn byte-cache
Adaptive compression:    Enabled
Adaptive compression index: 200
Max disk usage:          50%
(Max disk usage range should be between 5 and 80 percent of 126 GB)
```

Byte-Cache Configuration CLI Tables

As part of the byte-cache configuration CLI output, two tables are displayed:

- ❑ Global Information
- ❑ Per-Peer Data

Current Time (UTC)		Time of		Tot Size	Total	Tot Size		
Next Peer Ranking (UTC)		Allocabl		Rec Size	Alloc'd			
00:36:20 08/10/2009		03:05:00 08/10/2009		381 GB	4302 MB	4302 MB		
Peer ID	Traffic	Savings	Adj.	Rec.	Alloc.	Actual	Manual	Flags
			Gzip	Size	Size	Size	Size	
106080041	74 GB	73 GB	0 B	4302 MB	4302 MB	110 GB	0 B	

Viewing Byte-Cache Global Information

The first table has information that affects all caches, including the:

- ❑ current time
- ❑ time for the next scheduled (daily, at 3:05 AM local standard time) peer ranking
- ❑ total allocable disk space (converted from a percent into an actual size in SI units—20GB is 20,000,000,000 bytes)
- ❑ total recommended size of all dictionaries
- ❑ total allocated size of all dictionaries

Viewing Per-Peer Data

The second table has per-peer data, with one line for each peer (all ProxyClients are combined into a single line).

Note: All ProxyClients are shown on a single line. In this case it shows the total number of ProxyClients rather than the Peer ID. The corresponding statistics represent total overall client statistics for the traffic, savings, adjusted gzip, recommended size, allocated size, actual size, and manual size; the flags column displays an unbroken underline.

The following information is displayed:

- ❑ Peer ID—Peer ID of the peer ProxySG or the number of ProxyClients
- ❑ Traffic—Total uncompressed data over the last week
- ❑ Savings—Byte-cache savings during the last week
- ❑ Adj. Gzip—Adjusted gzip data (all the uncompressed data sent or received during the last week when byte caching was not being done)
- ❑ Rec. Size—Recommended size for this peer's dictionary
- ❑ Alloc. Size—Allocated size for this peer's dictionary
- ❑ Actual Size—Actual size for this peer's dictionary
- ❑ Manual Size—Manual size for this peer's dictionary
- ❑ Flags:
 - **N** indicates that the user chose not to do compression when sending data to this peer
 - **M** indicates that manual sizing is in effect for this dictionary
 - **A** indicates that the peer has advertised that it is using a manual size for its dictionary
 - **P** indicates that the dictionary is peer-limited. The peer has requested a smaller dictionary than allocated.

Chapter 33: WCCP Configuration

The Web Cache Communication Protocol (WCCP) is a Cisco-developed protocol that allows certain Cisco routers and switches to transparently redirect traffic to a cache engine such as a ProxySG appliance. This traffic redirection helps to improve response time and optimize network resource usage.

The ProxySG can be configured to participate in a WCCP scheme, in which WCCP-capable switches or routers collaborate with ProxySG appliances to form one or more groups that service requests from clients.

This section includes the following topics:

- ❑ ["WCCP on the ProxySG" on page 775](#)
- ❑ ["Prerequisites for Configuring WCCP on the ProxySG" on page 779](#)
- ❑ ["Configuring WCCP on the ProxySG" on page 780](#)
- ❑ ["Viewing WCCP Statistics and Service Group Status" on page 786](#)

WCCP on the ProxySG

In virtually in-path deployments, when the ProxySG is not in the physical path of clients and servers, a WCCP-capable router is used to redirect traffic to the ProxySG for transparent proxy services.

In a transparent proxy deployment the client is not aware that it is interacting with an intermediate proxy and not the OCS. The process works as follows:

1. The client sends a packet addressed for the OCS.
2. The WCCP-enabled router redirects the packet to the ProxySG.
3. The ProxySG determines what to do with it based on the transparent proxy services that have been configured for the traffic type. If it cannot service the request locally (for example by returning a page from its local cache), it sends a request to the specified OCS on behalf of the client.
4. The OCS response is routed (or redirected depending on the configuration) back to the ProxySG.
5. The ProxySG then forwards the response back to the client.

To implement this transparent redirection scheme, one or more ProxySG appliances and one or more routers/switches must form a service group.

The ProxySG offers VLAN Support for WCCP and allows you to redirect traffic from the router over physical or virtual interfaces. If you configure multiple virtual interfaces between the ProxySG and the WCCP-capable router, you can segregate WAN and LAN traffic on the same physical interface by enabling a VLAN trunk between the appliances. By default, VLAN trunking is enabled on the ProxySG. For information on configuring VLANs on the ProxySG, see ["About VLAN Configurations" on page 1233](#).

Service Groups

A service group unites one or more routers/switches with one or more ProxySG appliances in a transparent redirection scheme governed by a common set of rules. The service group members agree on these rules by announcing their specific capabilities and configuration to each other in WCCP protocol packets. When creating a service group on the ProxySG, you define the following:

- ❑ "Home Router Address" on page 776
- ❑ "Service Group Authentication" on page 776
- ❑ "Packet Forward and a Return Methods" on page 776
- ❑ "Router Affinity" on page 777
- ❑ "Assignment Types" on page 778
- ❑ "WCCP Load Balancing" on page 779

Home Router Address

In order to establish and maintain a service group, the ProxySG appliances and routers must be able to communicate. To specify the addresses of all the routers in a service group, you must choose one of the following methods:

- ❑ Unicast: Each ProxySG must be explicitly configured with the IP address of every router in the service group. You will need to reconfigure each ProxySG whenever you add or remove a router from the group.
- ❑ Multicast: The routers and ProxySG appliances in the service group communicate using a single IP address in the range of 224.0.0.0 to 239.255.255.255. To configure this, each ProxySG and each router in the group must be configured with the multicast IP address. If the WCCP routers and/or ProxySG appliances are more than one hop apart, IP multicast routing must also be enabled on the intervening routers.

Service Group Authentication

If you are using WCCP v2, you can secure a service group by configuring an MD5 authentication between the ProxySG appliances and the routers in the group. To configure authentication, you must define the same password on all routers and all ProxySG appliances in the service group.

When authentication is enabled, a ProxySG is not allowed to join the service group unless it knows the password.

Packet Forward and a Return Methods

The packet forward and return method for a service group defines how the router forwards packets to the ProxySG as well as how the ProxySG returns packets that it does not intercept because of the policy or services configured on it, back to the router.

Symantec recommends that all service groups configured on a router use the same forwarding and return methods.

The ProxySG supports the following forward/return methods:

- ❑ GRE Forwarding/GRE Return: With Generic Routing Encapsulation (GRE) forwarding, the router encapsulates the intercepted packet in an additional IP and GRE header that shows the router address as the source IP address and the address of the ProxySG as the destination IP address. When the ProxySG receives the packet, it strips the outside header and then determines how to process the request, either forwarding the request on to the OCS or servicing it locally.

When returning the redirected packet, the ProxySG encapsulates the packet with an IP and GRE header that bears the IP address of the ProxySG as the source and the router IP address as the destination.

- ❑ L2 Forwarding/L2 Return: With Layer 2 (L2) forwarding the router rewrites the destination MAC address of the intercepted packet to the MAC address of the ProxySG to which it is redirecting the packet. This method is faster than GRE forwarding because the forwarding is done at the hardware level and doesn't require encapsulating and decapsulating the packet at Layer 3. However, to use L2 forwarding, the ProxySG and the routers in the service group must all be on the same L2 broadcast domain (that is, there cannot be more than one hop between them).

When returning the redirected packet, the ProxySG rewrites the destination MAC address to that of the router.

To determine whether L2 forwarding is supported on your hardware platform, refer to your Cisco documentation. For a list of the Cisco platforms on which Symantec has tested L2 forwarding with the ProxySG, refer to the *WCCP Reference Guide*.

- ❑ L2Forwarding/GRE Return: With L2 forwarding the router rewrites the destination MAC address of the intercepted packet to the MAC address of the ProxySG to which it is redirecting the packet.

When returning the redirected packet, the ProxySG encapsulates the packet with an IP and GRE header that bears the IP address of the ProxySG as the source and the router IP address as the destination.

Note: The ProxySG does not support GRE forwarding and L2 packet return. If you configure this combination, the ProxySG will generate a capability mismatch error. To view the errors and warnings, click the **WCCP Status** button in the **Configuration > Network > WCCP** tab or use the CLI command `show wccp status`.

Router Affinity

By default, the ProxySG uses the configured return method to return bypassed traffic to the router that redirected it and uses regular routing table lookups to determine the next hop for intercepted traffic. With router affinity, the ProxySG also uses the configured return method to return intercepted client- and/or server-bound traffic to the WCCP router that redirected it, bypassing the routing table lookup. This is a useful feature if you have routing policies that may prevent your client- and/or server-bound traffic from reaching its destination and

simplifies the ProxySG configuration process by eliminating the need to replicate these policies on the ProxySG. It is also useful in configurations where you have multiple home routers or where your WCCP router is multiple hops away from the ProxySG because it ensures that the traffic is always returned to the same WCCP router that redirected it. Keep in mind, however, that enabling this feature unnecessarily when using GRE return does add additional CPU overhead on the router due to the need to decapsulate the GRE packets. In addition, the ProxySG and the router use a reduced maximum transmission unit (MTU) for GRE packets, which reduces the amount of data that can be transferred per packet.

Assignment Types

For every service group, you must configure the way the router determines the ProxySG to which to redirect a given packet, by setting an **assignment type** on the ProxySG. When the service group is formed, the ProxySG with the lowest IP address automatically becomes the *designated cache* (and if there is only one ProxySG in the service group, it is automatically the designated cache). The designated cache is responsible for communicating the assignment settings to the router, that is which ProxySG should be assigned a particular packet.

The ProxySG supports two assignment types:

- ❑ **Hash Assignment (Default):** With hash assignment, the designated cache assigns each ProxySG in the service group a portion of a 256-bucket hash table and communicates the assignment to the routers in the group. When the router receives a packet for redirection, it runs the hashing algorithm against one or more of the fields in the packet header to determine the hash value. It then compares the value to the hash assignment table to see which ProxySG is assigned to the corresponding bucket and then forwards the packet to that appliance. When you configure the service group on the ProxySG appliances, you specify which field(s)—destination IP address, destination port, source IP address, and/or source port—should be used to calculate the hash value.

In some cases, since all of the packets are hashed using the same fields and algorithm, it is possible that one of the caches in the group can become overloaded. For example, if you have a large proportion of traffic that is directed to the same server and you are using the destination IP address to run the hashing function, it is possible that the bulk of the traffic will be redirected to the same ProxySG. Therefore, you can configure an alternate field or group of fields to use to run the hashing algorithm. The router will then use this alternate hashing algorithm if the number of GRE packets or MAC addresses (depending on the forwarding method you're using) redirected to a given ProxySG exceeds a certain number.

For details on configuring a hash-weight value to adjust the proportion of the hash table that gets assigned to a ProxySG, see ["WCCP Load Balancing"](#) below.

- ❑ **Mask Assignment:** With mask assignment, each router in the service group has a table of masks and values that it uses to distribute traffic across the ProxySG appliances in the service group. When the router receives a packet, it performs a bitwise AND operation between the mask value and the field of the packet header that is designated in the ProxySG mask assignment configuration. It then compares the result against its list of values for each mask; each value is assigned to a specific ProxySG in the service group.

WCCP Load Balancing

Each ProxySG in the service group is assigned roughly an even percentage of the load by default, regardless of assignment type. If you would like to adjust or balance the load across multiple ProxySG appliances, you can assign a weight value to each ProxySG in the group. ProxySG appliances with higher weight values receive a larger portion of the redirected traffic.

For example, suppose you have assigned the following weight values: ProxySG1=100, ProxySG2=100, and ProxySG3=50 respectively. The total weight value is 250, and so ProxySG1 and ProxySG2 will each receive 2/5 of the traffic (100/250) and ProxySG3 will receive 1/5 of the traffic (50/250).

If a ProxySG becomes unavailable, the load will automatically be redistributed across the remaining ProxySG appliances in the service group.

Prerequisites for Configuring WCCP on the ProxySG

Before you configure WCCP on the ProxySG, you must complete the following tasks:

- ❑ **Plan your service groups:**
 - Decide which routers and which ProxySG appliances will work together in the redirection scheme.
 - Determine the WCCP capabilities that your router/switch supports. Refer to the documentation that came with your router for the specifics on your router/switch.
 - Decide what traffic you want to redirect. Do you want to redirect all traffic, or just a specific protocol or specific ports? Do you want to exclude certain hosts or traffic from redirection?
 - Decide what forwarding and return method you plan to use and make sure that all the routers in the service group support the chosen method(s).
 - Decide if you want to enable router affinity so that the ProxySG uses the chosen return method to return intercepted server- and/or client-bound traffic to the originating WCCP router as well as bypassed traffic.
 - Decide how the router will assign a specific redirected packet to a ProxySG. Make sure the router(s) in the service group support the assignment method you plan to use. If there is more than one ProxySG in the service group, decide whether you want to distribute traffic equally, or if you want to assign varying weights.

- ❑ Configure the routers. For information on the feature sets and the capabilities of your router and for instructions on how to configure WCCP on the router, refer to the router documentation. For sample router WCCP configurations, refer to the *WCCP Reference Guide*.

Configuring WCCP on the ProxySG

You must configure the required WCCP settings on the participating routers before proceeding with this section.

Use the procedures in this section to perform the following tasks:

- ❑ ["Creating the WCCP Configuration on the ProxySG"](#) on page 780.
- ❑ ["Modifying the WCCP Configuration"](#) on page 785.
- ❑ ["Disabling WCCP"](#) on page 786.

Creating the WCCP Configuration on the ProxySG

You must create a WCCP configuration file on the ProxySG that contains the WCCP settings specific to the ProxySG. When installed, these configuration settings enable the ProxySG to collaborate with the WCCP-capable router or switch.

You can create the WCCP configuration file in three ways:

- ❑ Using the user interface in Management Console. This option provides a graphical interface that prompts you to select from the options on-screen and enter values as appropriate. For instructions, see ["Configuring WCCP from the Management Console"](#) on page 780.
- ❑ Using a text editor. This option allows you to create and install a text file on:
 - a remote machine and access the URL through the Management Console.
 - a file locally on the system from which you run the Management Console.
 - the text editor in the Management Console. The Management Console provides a text editor that can be used to create the configuration file. You can copy and paste the contents of an existing configuration file or you can enter new text.

For descriptions of the values in the configuration file, refer to the *WCCP Reference Guide*. For instructions on installing the settings, see ["Configuring WCCP Settings Using the Text Editor"](#) on page 784.


- ❑ Using the `inline wccp-settings eof_marker` CLI command to type the WCCP configuration using the terminal. For more information, refer to the *WCCP Reference Guide*.

Configuring WCCP from the Management Console

The easy-to-use interface allows you to configure the WCCP settings on the ProxySG. For a description of the configuration options, see ["WCCP on the ProxySG"](#) on page 775.

To create the WCCP configuration using the Management Console:


1. Select the **Configuration > Network > WCCP** tab.



2. Select **Enable WCCP**.
3. Select the **WCCP Version**. Unless you are creating a web-cache service group, you must use version 2.0.
4. Click **Apply**.

Note: If you select version 1.0, you can only configure a single web-cache service group. The web-cache service group is a well-defined service group that intercepts all TCP traffic on destination port 80. When configuring a web-cache service group, you must select an interface to which apply the service group and define a single home router. You can optionally enable router affinity. See "[Router Affinity](#)" on page 777 for more information on this setting.

5. To create a service group, click **New**. The **New Service** dialog displays.



6. Define the service group and apply it to an interface:
 - a. Enter a **Service Group** number. The service group number must be a unique identifier in the range of 0 to 255 inclusive.
 - b. (Optional) Specify the service group **Priority** in the range of 0-255. When multiple service groups that are redirecting the same traffic (for example HTTP on port 80) are assigned to a common router interface, the priority defines the order in which the router evaluates the service groups.
 - c. (Optional) Set a **Password** to configure MD5 authentication for added security. The password can include 1 to 8 characters. When authentication is enabled on the router, the ProxySG must be configured with the same password to join the service group.

7. Apply the service group to one or more physical or a virtual interfaces.
 - a. Select a value from the **Interface** drop-down menu. Virtual interfaces are depicted as *adapter:interface.vlan id*, for example, 0:0.3.
 - b. (Optional) Enter the **Weight** value for this interface. This value determines the proportion of traffic that the router redirects to this interface. The weight value can range from 1 to 255 inclusive. Use this field only if you are redirecting traffic in the same service group to multiple interfaces or to multiple ProxySG appliances and you want to allocate the percentage of traffic redirected to each ProxySG and/or interface in the service group.
 - c. To add additional interfaces on this ProxySG appliance to the service group, click **Add Interface** and then repeat steps **a** and **b**.

8. Define the traffic you want to redirect (ports and protocols).
 - a. (Optional) If you want to redirect specific ports instead of redirecting **All** traffic (the default), select a value from the **Redirect on** drop-down list. You can choose from **Source**, **Destination**, or **All**.
 - b. Select the **Protocol** to redirect — **TCP** or **UDP**.
 - c. Specify the **Ports to redirect**. If you selected **Source** or **Destination** from the **Redirect on** drop-down list, you must select the applicable options and/or specify ports in the **Other** field. You can specify up to 8 ports to redirect for each service group. If you want to redirect more than 8 ports, you must create more than one service group.

9. Define how the router and the ProxySG handle packet forwarding and return:
 - a. Select a **Forwarding Type** — Generic Routing Encapsulation (**GRE**) or Layer 2 forwarding (**L2**). For a description of these options, see ["Packet Forward and a Return Methods"](#) on page 776.

- b. Select a **Returning Type**. Only applicable if you select L2 forwarding. For the GRE forwarding method, the ProxySG only supports GRE return.
 - c. (Optional) If you want to ensure that intercepted traffic is always routed through the WCCP router that redirected it, select a **Router affinity** value:
 - **Client** indicates that the ProxySG will return client-bound traffic to the originating WCCP router using the configured **Returning Type**.
 - **Server** indicates that the ProxySG will return server-bound traffic to the originating WCCP router using the configured **Returning Type**.
 - **Both** indicates that the ProxySG will return both client- and server-bound traffic to the originating WCCP router using the configured **Returning Type**.
 - **<None>** (the default) indicates that the ProxySG will use regular routing table lookups rather than the configured **Returning Type** to route the client- and server-bound traffic that it intercepts.
10. Add the home router address. Specify individual unicast or a single multicast address for the router(s) in the service group:
- If you want to use multicast addressing, select **Multicast Home Router** and enter the **Group Address** and optionally a **Multicast TTL** value (default =1).
 - If you want to use unicast addresses, select **Individual Home Router Address**. For each router in the service group, click **Add**, enter the **Home Router Address** and click **OK**. The home router address that you use for a service group on the ProxySG should be consistent with the IP address (virtual or physical) over which the ProxySG communicates with the router.
11. Select an **Assignment Type**. The assignment type instructs the router how to distribute redirected traffic using the information in the packet header. You can select a different assignment method for each service group configured on the same ProxySG.

Assignment Type: ☒ Hash ☐ Mask

Primary Hash : ☒ Source IP ☐ Source Port ☐ Destination IP ☐ Destination Port

Alternate Hash: ☐ Source IP ☐ Source Port ☐ Destination IP ☐ Destination Port

- If you select the **Hash** assignment type (the default), you can select one or more fields to use as the **Primary Hash**. Additionally, you can optionally select one or more fields to use as the **Alternate Hash**. The alternate hashing function is used to distribute traffic when a particular ProxySG exceeds a given number of redirected packets.

Assignment Type: ☐ Hash ☒ Mask

Mask Scheme: ☒ Source IP ☐ Source Port ☐ Destination IP ☐ Destination Port

Mask Value: Enter mask as hex value prefixed by 0x or as decimal.

- If you select the **Mask** assignment type, select which field in the packet header to use to run the mask function. Enter a **Mask Value** in either decimal or, when prefixed by 0x, a hexadecimal value. The default value for this field is 0x3f. The following Cisco Web page describes the Mask Value in detail: http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11-629052.html.

1. Click **OK** to save the service group settings. If you want to add another service group, repeat Steps 5 through 1.
2. To save the WCCP settings, click **Apply**.

Configuring WCCP Settings Using the Text Editor

Whether you opt to use the text editor in the Management Console, text editor on the local system, or you plan to install the configuration from a remote file, use the instructions below to install the WCCP settings on the ProxySG.

To install the configuration file:

1. Select the **Configuration > Network > WCCP** tab.
2. Select **Enable WCCP**.

The screenshot shows the 'Install WCCP Settings' panel in the management console. It includes a section 'Install WCCP Settings from:' with a dropdown menu currently showing 'Remote URL'. Other options in the dropdown are 'Local File' and 'Text Editor'. An 'Install' button is visible to the right. Below this, the 'View WCCP Settings' section contains two buttons: 'WCCP Settings' and 'WCCP Source', each with a corresponding 'View' link.

3. In the **Install WCCP Settings** panel, select the location of the configuration file: a remote URL, a local file, or use the text editor on the system.
4. Click **Install**.
If you selected **Remote URL** or **Local File**, a dialog opens that allows you to enter the complete path, and the file is retrieved. If you selected **Text Editor**, the text editor displays with the current settings. You can copy and paste the contents of an existing configuration file or you can enter new text and click **Install** when finished.

The following shows an example WCCP configuration:

```
wccp enable
wccp version 2
service-group 9
  forwarding-type L2
  returning-type GRE
  router-affinity both
  assignment-type mask
  mask-scheme source-port
  priority 1
  protocol 6
```



```

service-flags ports-defined
ports 80 21 1755 554 0 0 0 0
interface 0:0
home-router 10.16.18.2
end

```

For descriptions of the settings in the configuration file, refer to the *WCCP Reference Guide*.


- (Optional): View the WCCP settings that are currently on the system or view the text file with the current settings by clicking **WCCP Settings** or **WCCP Source**.
- Click **Apply** to save the changes.

Modifying the WCCP Configuration

The following sections describe how to modify or delete a service group. For instructions on adding a service group, see "[Configuring WCCP from the Management Console](#)" on page 780.

To edit a service group:

- Select the **Configuration > Network > WCCP** tab.



The screenshot shows the 'WCCP Configuration' window. At the top, there are radio buttons for 'WCCP Version : 1.0' and '2.0', with '2.0' selected. Below this is a table with four columns: 'Groups', 'Home Router', 'Ports', and 'State'. There are two rows of data. The first row has '3' in the Groups column, '1.1.1.1, 2.2.2.2' in the Home Router column, 'TCP Source Ports: 80 (HTTP), 443 (HTTPS), 139 (CIFS)...' in the Ports column, and 'Negotiating me...' in the State column. The second row has '4' in the Groups column, '224.0.0.0' in the Home Router column, 'TCP Source Ports: 80 (HTTP), 443 (HTTPS)' in the Ports column, and 'Negotiating me...' in the State column. Below the table are four buttons: 'New', 'Edit', 'Delete', and 'Refresh State'.

Groups	Home Router	Ports	State
3	1.1.1.1, 2.2.2.2	TCP Source Ports: 80 (HTTP), 443 (HTTPS), 139 (CIFS)...	Negotiating me...
4	224.0.0.0	TCP Source Ports: 80 (HTTP), 443 (HTTPS)	Negotiating me...

- Select the service group to modify.
- Click **Edit**. The **Edit Service** dialog displays.
- Perform the changes. You can edit any value except for the service group number.
- Click **OK**.
- Click **Apply** to save your changes.

To delete a service group:

- Select the **Configuration > Network > WCCP** tab.
- Select the service group that you want to delete.
- Click **Delete**. The service group is deleted; you are not prompted for confirmation.
- Click **Apply**.

Disabling WCCP

To exclude a ProxySG from receiving traffic or from participating in any of the services groups configured on it, you can disable WCCP on the ProxySG. Disabling WCCP does not delete the WCCP configuration settings, it places them out-of-service until WCCP is re-enabled on the ProxySG.

To disable WCCP on the ProxySG:

1. Select the **Configuration > Network > WCCP** tab.
2. Clear the **Enable WCCP** check box. When WCCP is disabled, the previous WCCP statistics are cleared.
3. Click **Apply** to save your changes.

Viewing WCCP Statistics and Service Group Status

After you install the WCCP configuration, the WCCP routers and ProxySG appliances in the defined service groups begin negotiating the capabilities that you have configured. You can monitor the statistics for the configured service groups either from the Management Console or from the CLI of the ProxySG.

To view WCCP statistics:

Select **Statistics > Network > WCCP**. The top of the page displays whether WCCP is enabled. If WCCP is disabled, no statistics are displayed.

The screenshot shows the WCCP configuration page. At the top, it says "WCCP is enabled." Below this, there's a "WCCP Status" section. It includes a "Last Refresh" timestamp of "Thu Mar 20 12:19:26 PDT 2008" and a "Refresh WCCP Statistics" button. Below the button, it shows "GRE Redirected Packets: 2,004,678" and "Layer-2 Redirected Packets: 0". The main part of the screenshot is a table with columns: "Services Groups", "State", "Here I Am Sent", "I See You Received", and "Redirect Assign Sent". There are two rows of data, both for "Group: 19" (Note: the image shows Group 9 and Group 19, but the text in the image is Group 9 and Group 19). Both groups are in a "Ready" state. The "Here I Am Sent" and "I See You Received" values are both 867. The "Redirect Assign Sent" value is 1. Each group entry also shows "Cache: 192.168.20.2" and "Router: 192.168.20.1".

Services Groups	State	Here I Am Sent	I See You Received	Redirect Assign Sent
Group: 9 Cache: 192.168.20.2 Router: 192.168.20.1	Ready	867	867	1
Group: 19 Cache: 192.168.20.2 Router: 192.168.20.1	Ready	867	867	1

If WCCP is disabled, the following statistics are displayed:

Statistic	Description
Last Refresh	The date and time the displayed statistics were last refreshed. Click Refresh WCCP Statistics to refresh them now.
GRE Redirected Packets	The number of packets that have been redirected using GRE forwarding.
Layer-2 Redirected Packets	The number of packets that have been redirected using L2 forwarding.

Statistic	Description
Services Groups	Lists the service groups that have been configured on this ProxySG. If the group has successfully formed, you can click the arrow next to the group to see a list of the caches (ProxySG appliances) and routers that have joined the group.
State	Shows the service group state. See Table 17–1 for a description of each state.
Here I Am Sent	The number of HERE_I_AM messages that this ProxySG has sent to the routers in the group.
I See You Received	The number of I_SEE_YOU messages that this ProxySG has received from the routers in the group.
Redirect Assign Sent	The number of REDIRECT_ASSIGN messages that this ProxySG has sent to the routers in the group. The REDIRECT_ASSIGN message contains the hash table or mask values table that the router will use to determine which ProxySG to redirect packets to. Only the designated cache—the cache with the lowest IP address—sends REDIRECT_ASSIGN messages.

Monitoring the Service Group States

The ProxySG maintains state information on the configured service groups. The state of a service group helps you monitor whether the service group was configured properly and on how it is functioning.

To view the state of the service groups you have configured, see **Statistics > Network > WCCP**

Table 17–1 lists and describes each service group state.

Table 33–1 WCCP Service Group States

State	Description
Assignment mismatch	The router does not support the assignment type (hash or mask) that is configured for the service group.
Bad router id	The home-router specified in the service group configuration does not match the actual router ID.
Bad router view	The list of ProxySG appliances in the service group does not match.
Capability mismatch	The WCCP configuration includes capabilities that the router does not support.
Initializing	WCCP was just enabled and the ProxySG is getting ready to send out its first HERE_I_AM message.
Interface link is down	The ProxySG cannot send the HERE_I_AM message because the interface link is down.
Negotiating assignment	The ProxySG received the I_SEE_YOU message from the router but has not yet negotiated the service group capabilities.
Negotiating membership	The ProxySG sent the HERE_I_AM message and is waiting for an I_SEE_YOU message from the router.

Table 33–1 WCCP Service Group States

Packet forwarding mismatch	The router does not support the forwarding method (GRE or L2) that is configured for the service group.
Packet return mismatch	The router does not support the return method (GRE or L2) that is configured for the service group. Note that on the ProxySG, the return method is always the same as the forwarding method.
Ready	The service group formed successfully and the ProxySG sent the REDIRECT_ASSIGN message to the router with the hash or mask values table.
Service group mismatch	The router and the ProxySG have a mismatch in port, protocol, priority, and/or other service flags.
Security mismatch	The service group passwords on the router and the ProxySG do not match.

Chapter 34: TCP/IP Configuration

This chapter describes the TCP/IP configuration options, which enhance the performance and security of the ProxySG. Except for IP Forwarding, these commands are only available through the CLI.

Topics in this Chapter

The following topics are discussed in this chapter:

- ❑ ["About the Options"](#) on page 789
- ❑ ["RFC-1323"](#) on page 790
- ❑ ["TCP NewReno"](#) on page 790
- ❑ ["ICMP Broadcast Echo Support"](#) on page 790
- ❑ ["ICMP Timestamp Echo Support"](#) on page 790
- ❑ ["To configure network tunneling settings:"](#) on page 791
- ❑ ["PMTU Discovery"](#) on page 791
- ❑ ["TCP Time Wait"](#) on page 793
- ❑ ["TCP Loss Recovery Mode"](#) on page 793
- ❑ ["Viewing the TCP/IP Configuration"](#) on page 794

About the Options

- ❑ **RFC-1323:** Enabling RFC-1323 support enhances the high-bandwidth and long-delay operation of the ProxySG appliances over very high-speed paths, ideal for satellite environments.
- ❑ **TCP NewReno:** Enabling TCP NewReno support improves the fast recovery of the appliances.
- ❑ **ICMP Broadcast Echo:** Disabling the response to these messages can limit security risks and prevent an attacker from creating a distributed denial of service (DDoS) to legitimate traffic.
- ❑ **ICMP Timestamp Echo:** Disabling the response to these messages can prevent an attacker from being able to reverse engineer some details of your network infrastructure.
- ❑ **TCP Window Size:** Configures the amount of unacknowledged TCP data that the ProxySG can receive before sending an acknowledgement.
- ❑ **PMTU Discovery:** Enabling PMTU Discovery prevents packets from being unable to reach their destination because they are too large.

To view the TCP/IP configuration, see ["TCP Loss Recovery Mode"](#) on page 793.

RFC-1323

The RFC-1323 TCP/IP option enables the ProxySG to use a set of extensions to TCP designed to provide efficient operation over large bandwidth-delay-product paths and reliable operation over very high-speed paths, including satellite environments. RFC-1323 support can be configured through the CLI and is enabled by default.

To enable or disable RFC-1323 support:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) tcp-ip rfc-1323 {enable | disable}
```

TCP NewReno

NewReno is a modification of the Reno algorithm. TCP NewReno improves TCP performance during fast retransmit and fast recovery when multiple packets are dropped from a single window of data. TCP NewReno support is enabled by default.

To enable or disable TCP NewReno support:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) tcp-ip tcp-newreno {enable | disable}
```

ICMP Broadcast Echo Support

Disabling the ICMP broadcast echo command can prevent the ProxySG from participating in a Smurf Attack. A Smurf attack is a type of Denial-of-Service (DoS) attack, where the attacker sends an ICMP echo request packet to an IP broadcast address. This is the same type of packet sent in the `ping` command, but the destination IP is broadcast instead of unicast. If all the hosts on the network send echo reply packets to the ICMP echo request packets that were sent to the broadcast address, the network is jammed with ICMP echo reply packets, making the network unusable. By disabling ICMP broadcast echo response, the ProxySG does not participate in the Smurf Attack.

This setting is disabled by default.

To enable or disable ICMP broadcast echo support:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) tcp-ip icmp-bcast-echo {enable | disable}
```

For more information on preventing DDoS attacks, see [Chapter 70: "Preventing Denial of Service Attacks"](#) on page 1279.

ICMP Timestamp Echo Support

By disabling the ICMP timestamp echo commands, you can prevent an attacker from being able to reverse engineer some details of your network infrastructure.

For example, disabling the ICMP timestamp echo commands prevents an attack that occurs when the ProxySG responds to an ICMP timestamp request by accurately determining the target's clock state, allowing an attacker to more effectively attack certain time-based pseudo-random number generators (PRNGs) and the authentication systems on which they rely.

This setting is disabled by default.

To enable or disable ICMP Timestamp echo support:

At the (config) command prompt, enter the following command:

```
SGOS#(config) tcp-ip icmp-timestamp-echo {enable | disable}
```

PMTU Discovery

Path MTU (PMTU) discovery is a technique used to determine the maximum transmission unit (MTU) size on the network path between two IP hosts to avoid IP fragmentation.

A ProxySG that is not running PMTU might send packets larger than that allowed by the path, resulting in packet fragmentation at intermediate routers. Packet fragmentation affects performance and can cause packet discards in routers that are temporarily overtaxed.

A ProxySG configured to use PMTU sets the `Do-Not-Fragment` bit in the IP header when transmitting packets. If fragmentation becomes necessary before the packets arrive at the second ProxySG, a router along the path discards the packets and returns an `ICMP Host Unreachable` error message, with the error condition of `Needs-Fragmentation`, to the original ProxySG appliance. The first appliance then reduces the PMTU size and re-transmits the transmissions.

The discovery period temporarily ends when the ProxySG estimates the PMTU is low enough that its packets can be delivered without fragmentation or when the ProxySG stops setting the `Do-Not-Fragment` bit.

Following discovery and rediscovery, the size of the packets that are transferred between the two communicating nodes dynamically adjust to a size allowable by the path, which might contain multiple segments of various types of physical networks.

Note: PMTU is disabled by default.

To configure PMTU discovery:

At the (config) command prompt:

```
SGOS#(config) tcp-ip pmtu-discovery {enable | disable}
```

To configure network tunneling settings:

1. Select **Configuration > ADN > Tunneling > Network**.

The screenshot shows the 'Network' tab in the SGOS configuration interface. It contains two sections: 'Reflect Client IP' and 'TCP Settings'.

Reflect Client IP

When an inbound ADN tunnel connection requests Client IP reflection:

- ☐ Reject the request
- ☐ Allow the request and reflect the client IP
- ☒ Allow the request but connect using a local IP

TCP Settings

ADN Tunnel TCP Window Size:

- ☒ Automatically adjusted
- ☐ Manual override:

2. Determine the behavior of the concentrator proxy when a branch proxy requests client IP reflection (sending the client's IP address instead of the ProxySG IP address to the upstream server).

This setting is based on whether the concentrator was installed inline. If the concentrator proxy is inline and can do IP reflection, you can allow client IP address reflection requests from clients. If not, set this option to either **Reject the Request** or **Allow the request but connect using a local IP** to accept the requests but ignore the client IP address and use a local IP address.

3. *TCP window size* is the number of bytes that can be buffered on a system before the sending host must wait for an acknowledgement from the receiving host.

The TCP window size for ADN optimization tunnel connections is set and updated automatically, based on current network conditions and on the receiving host's acknowledgement. In most situations, the **TCP Settings** option should be left as **Automatically adjusted**.

Only use the **Manual override** setting if your network environment has intervening network equipment that makes the delay appear lower than it actually is. These environments are sometimes found on satellite links that have high bandwidth and high delay requirements. In this case, the automatically adjusted window size would be smaller than optimal.

The configurable range is between 8 Kb and 4 MB (8192 to 4194304), depending on your bandwidth and the round-trip delay. Setting sizes below 64Kb are not recommended.

Note: If you know the bandwidth and round-trip delay, you can compute the value to use as, roughly, $2 * \text{bandwidth} * \text{delay}$. For example, if the bandwidth of the link is 8 Mbits/sec and the round-trip delay is 0.75 seconds:

$$\text{window} = 2 * 8 \text{ Mbits/sec} * 0.75 \text{ sec} = 12 \text{ Mbits} = 1.5 \text{ Mbytes}$$

The setting in this example would be 1500000 bytes. This number goes up as either bandwidth or delay increases, and goes down as they decrease.

You can decrease or increase the window size based on the calculation; however, decreasing the window size below 64Kb is not recommended..

The window-size setting is a maximum value; the normal TCP/IP behaviors adjust downward as necessary. Setting the window size to a lower value might result in an artificially low throughput.

4. Click **Apply** to commit the changes to the ProxySG.

TCP Time Wait

When a TCP connection is closed (such as when a user enters *quit* for an FTP session), the TCP connection remains in the `TIME_WAIT` state for twice the Maximum Segment Lifetime (MSL) before completely removing the connection control block.

The `TIME_WAIT` state allows an end point (one end of the connection) to remove remnant packets from the old connection, eliminating the situation where packets from a previous connection are accepted as valid packets in a new connection.

The MSL defines how long a packet can remain in transit in the network. The value of MSL is not standardized; the default value is assigned according to the specific implementation.

To change the MSL value, enter the following commands at the (config) command prompt:

```
SGOS#(config) tcp-ip tcp-2msl seconds
```

where *seconds* is the length of time you chose for the 2MSL value. Valid values are 1 to 16380 inclusive.

TCP Loss Recovery Mode

The TCP loss recovery mode algorithm helps recover throughput efficiently after random packet losses occur over your network, such as across wireless and satellite paths. It also addresses performance problems due to a single packet loss during a large transfer over long delay pipes, such as transcontinental or transoceanic pipes.

The TCP loss recovery mode is set to *Normal* by default. Symantec recommends that you consider non-normal loss modes — enhanced or aggressive, only when you experience packet losses of 0.5% or greater. For quickly estimating packet losses between two endpoints in your network, you can run a continuous, non-flooding, *ping* for a couple minutes and record the reported loss.

Symantec does not recommend modifying the loss recovery mode to *Aggressive* unless your network has demonstrated an improvement in the *Enhanced* mode. *Aggressive* mode may not provide further improvement, and in some instances it could worsen network performance. For additional information and guidance, contact Symantec Technical Support.

If you reconfigure the TCP loss recovery mode, you must configure the TCP window size that is appropriate for the link. A good rule is to set the window size to bandwidth times round-trip delay. For example, a 1.544Mbps link with a 100ms round-trip time would have a window size of 19,300 bytes.

The bandwidth and round-trip times can be determined from link characteristics (such as from the ISP) or observations (such as ping usage).

To set the TCP loss recovery algorithm to a non-normal mode:

```
SGOS#(config) tcp-ip tcp-loss-recovery-mode {enhanced | aggressive}
```

To reset the TCP loss recovery algorithm to the normal mode:

```
SGOS#(config) tcp-ip tcp-loss-recovery-mode {normal}
```

Viewing the TCP/IP Configuration

To view the TCP/IP configuration:

```
SGOS#(config) show tcp-ip
RFC-1323 support:          enabled
TCP Newreno support:       disabled
IP forwarding:             disabled
ICMP bcast echo response:  disabled
ICMP timestamp echo response: disabled
Path MTU Discovery:        disabled
TCP 2MSL timeout:          120 seconds
TCP window size:           65535 bytes
TCP Loss Recovery Mode:    Normal
```

Chapter 35: Routing on the ProxySG

This chapter explains how the ProxySG delivers packets and describes the features you can use to optimize packet delivery.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ "Basic Traffic Routing" on page 795
- ❑ "Distributing Traffic Through Multiple Default Gateways" on page 796
- ❑ "Configuring IP Forwarding" on page 798
- ❑ "Outbound Routing" on page 798
- ❑ "DNS Verification" on page 805

Basic Traffic Routing

Because it does not participate in a network routing protocol, the ProxySG must be configured to reach clients and servers. To reach devices outside the network, you must configure a primary packet delivery path. This path is known as the default route (or the *default gateway*) and is configured during initial setup when you specify a default gateway. The ProxySG sends all traffic to the default gateway unless another route is specified. These alternate routes are called *static routes* and they list the IP addresses of other gateways that can be used to reach clients and servers in other parts of the network. Static routes are discussed in "About Static Routes" on page 800.




Figure 35–1 Network example of default and static route

The ProxySG can be configured to distribute traffic through multiple default gateways, as explained in the next section.

Distributing Traffic Through Multiple Default Gateways

You can distribute traffic originating at the ProxySG through *multiple* default gateways and fine tune how the traffic is distributed. This feature works with any routing protocol.

By using multiple gateways, an administrator can assign a number of available gateways into a preference group and configure the load distribution to the gateways within the group. Multiple preference groups are supported.

In a mixed IPv4/IPv6 environment, you need to define two default gateways: one for IPv4 and one for IPv6.

The specified gateway applies to all network adapters in the system.

Note: Load balancing through multiple gateways is independent from the per-interface load balancing the ProxySG automatically does when more than one network interface is installed.

ProxySG Specifics

Which default gateway the ProxySG uses at a given time is determined by the preference group configuration assigned by the administrator. A ProxySG can have from 1 to 10 preference groups. A group can contain multiple gateways or only a single gateway.

Each gateway within a group can be assigned a relative weight value from 1 to 100. The weight determines how much bandwidth a gateway is given relative to the other gateways in the same group. For example, in a group with two gateways, assigning both gateways the same weight, whether 1 or 100, results in the same traffic distribution pattern. Alternatively, assigning one gateway a value of 10 and the other gateway a value of 20 results in the ProxySG sending approximately twice the traffic to the gateway with a weight value of 20.

If there is only one gateway, it automatically has a weight of 100.

All gateways in the lowest preference group are considered to be active until one of them becomes unreachable and is dropped from the active gateway list. Any remaining gateways within the group continue to be used. If all gateways in the lowest preference group become unreachable, the gateways in the next lowest preference group become the active gateways (unless a gateway in a lower preference group becomes reachable again).

Note: Unreachable means that ICMP or other requests have failed, either because the system is down or because the network path failed.


Switching to a Secondary Default Gateway

When a gateway becomes unreachable, the networking code detects the unreachable gateway in 20 seconds, and the switchover takes place immediately if a secondary gateway is configured.


For more information, see ["Distributing Traffic Through Multiple Default Gateways"](#) on page 796.

To configure multiple gateway load balancing:

1. Select the **Configuration > Network > Routing > Gateways** tab.



2. Click **New**. The Add List Item dialog displays.



3. Configure the gateway options:
 - a. In the **Gateway** field, enter the gateway IP address (IPv4 or IPv6).
 - b. From the **Group** drop-down list, select the preference group for this gateway.
 - c. In the **Weight** field, enter the relative weight within the preference group.
 - d. Click **OK** to close the dialog.
4. Repeat steps 2 to 4 until IP addresses, groups, and weights have been defined for all of your gateways.
5. Click **Apply**.

Configuring IP Forwarding

IP Forwarding is a special type of transparent proxy. The ProxySG is configured to act as a gateway and is configured so that if a packet is addressed to the ProxySG adapter, but not its IP address, the packet is forwarded toward the final destination. If IP forwarding is disabled, the packet is rejected as being mis-addressed.

By default, IP forwarding is disabled to maintain a secure network.

Important: When IP forwarding is enabled, be aware that all ProxySG ports are open and all the traffic coming through them is not subjected to policy, with the exception of the ports that have explicitly defined through the **Configuration > Services > Proxy Services** tab.

To enable IP forwarding:

1. Select **Configuration > Network > Routing > Gateways**.
2. Select the **Enable IP forwarding** check box at the bottom of the pane.
3. Click **OK**; click **Apply**.

How Routes are Determined

Typically, the ProxySG uses the routing table to determine which interface to send the outbound packets to. When a packet is received, the ProxySG does a routing lookup to see if it can determine the correct route, either by using a static route or, if one is not defined, by sending it over the default route.

However, the routing lookup might be bypassed depending on how the ProxySG is deployed and configured, as explained in the next section.

Routing in Transparent Deployments

This section describes the mechanisms the ProxySG uses to route packets.

Outbound Routing

By default, the ProxySG sends outbound traffic to the default gateway unless one of the following is used (in order of precedence):

- ❑ The Trust Destination MAC feature, which is used when the ProxySG is in transparent bridging mode (unless certain other conditions are true—see ["About Trust Destination MAC"](#) on page 799).

- ❑ A static route, if one is defined.

For more information, see ["About Static Routes"](#) on page 800.

- ❑ The outbound Return-to-Sender (RTS) feature.

For more information, see ["Using Return-to-Sender \(RTS\)"](#) on page 801.

- ❑ An interface route, if the device is on the same subnet as the ProxySG.

The appliance automatically adds an interface route to the routing table for hosts on the same subnet as the ProxySG interface. The interface route maps the subnet to the interface. The ProxySG can then do an ARP lookup for those hosts and send the packets directly to the client's MAC address.

Inbound Routing

By default, the ProxySG sends inbound traffic to the default gateway unless one of the following is used (in order of precedence):

- ❑ A static route, if one is defined.

For more information, see ["About Static Routes"](#) on page 800.

- ❑ The inbound RTS feature. Inbound RTS is enabled by default.

For more information, see ["Using Return-to-Sender \(RTS\)"](#) on page 801.

- ❑ An interface route, if the device is on the same subnet as the ProxySG.

About Trust Destination MAC

Starting with SGOS 5.1.4.x, when the ProxySG is in transparent bridging mode (in-path), it "trusts" the destination MAC address of the first client SYN packet and does not consult its routing table. The ProxySG notes the destination MAC address and outgoing interface specified in the frame and passes that information to the software process initiating the server connection, thus avoiding a routing lookup on the ProxySG. This feature is called Trust Destination MAC. It is enabled by default when the ProxySG is in transparent bridging mode and cannot be disabled.

Trust Destination MAC eliminates the need to create static routes and circumvents any routing issues encountered when the information in a packet is not sufficient for the ProxySG to make a routing decision.

Note: Trust Destination MAC uses only the first client SYN packet to determine the MAC address and outgoing interface and continues to use this information even if the destination MAC address is not responding. To work around this limitation, enable outbound RTS, as described in ["Using Return-to-Sender \(RTS\)"](#) on page 801.

Overriding Trust Destination MAC

Unlike RTS, non-default static routes cannot override Trust Destination MAC. However, Trust Destination MAC behavior can be overridden if any of the following conditions are true:

1. The result of the DNS lookup does not match the destination IP address. If Trust Destination IP is enabled, the DNS lookup is bypassed, and the IP address should always match.

2. A policy rule does one of the following:
 - Specifies a forwarding host or SOCKS gateway
 - Rewrites the server URL in a way that causes the server connection to be forwarded to a different host

About Static Routes

Static routes define alternate gateways that the ProxySG can send packets to. A static route is a manually-configured route that specifies a destination network or device and the specific router that should be used to reach it. See "[Defining Static Routes](#)".

Note: For transparent bridge deployments, Trust Destination MAC overrides any static routes.

Defining Static Routes

You define static routes in a routing table. The routing table is a text file containing a list of static routes made up of destination IP addresses (IPv4 or IPv6), subnet masks (for IPv4) or prefix lengths (for IPv6), and gateway IP addresses (IPv4 or IPv6). You are limited to 10,000 entries in the static routes table. The following is a sample routing table:

```
10.25.36.0    255.255.255.0    10.25.36.1
10.25.37.0    255.255.255.0    10.25.37.1
10.25.38.0    255.255.255.0    10.25.38.1
2001::/64      fe80::2%0:1
```

Note that a routing table can contain a combination of IPv4 and IPv6 entries, but the gateway for each destination must be on the appropriate network type. For example, an IPv6 destination must use an IPv6 gateway.

When a routing table is installed, all requested URLs are compared to the list and routed based on the best match.

You can install the routing table several ways.

- ❑ Using the Text Editor, which allows you to enter settings (or copy and paste the contents of an already-created file) directly onto the appliance.
- ❑ Creating a local file on your local system; the ProxySG can browse to the file and install it. See "[Installing a Routing Table](#)" on page 801.
- ❑ Using a remote URL, where you place an already-created file on an FTP or HTTP server to be downloaded to the ProxySG. Use the `static-routes path` command to set the path and the `load static-route-table` command to load the new routing table.
- ❑ Using the CLI `inline static-route-table` command, which allows you to paste a static route table into the ProxySG.

Note: If you upgrade to SGOS 5.x from SGOS 4.x, entries from the central and local bypass lists are converted to static route entries in the static route table. The converted static route entries are appended after the existing static route entries. Duplicate static route entries are silently ignored.

All traffic leaving the ProxySG is affected by the static route entries created from the SGOS 4.x bypass lists.

Installing a Routing Table

To install a routing table:

1. Select **Configuration > Network > Routing > Routing**.
2. From the drop-down list, select the method used to install the routing table; click **Install**.
 - **Remote URL:**

Enter the fully-qualified URL, including the filename, where the routing table is located. To view the file before installing it, click **View**. Click **Install**. To view the installation results, click **Results**; close the window when you are finished. Click **OK**.
 - **Local File:**

Click **Browse** to bring up the Local File Browse window. Browse for the file on the local system. Open it and click **Install**. When the installation is complete, a results window opens. View the results and close the window.
 - **Text Editor:**

The current configuration is displayed in installable list format. You can customize it or delete it and create your own. Click **Install**. When the installation is complete, a results window opens. View the results, close this window, and click **Close**.
3. Click **Apply**.

Note: If you use URL host rewrite functionality in your policies, mismatches can occur between the client-provided IP address and the resolved, rewritten hostname. In these cases, a routing lookup is performed and an interface route, static route, or default route is used.

Using Return-to-Sender (RTS)

As stated previously, the ProxySG does a routing lookup to see if it can determine the correct route for a packet, either by using a static route or, if one is not defined, by sending it over the default route. However, using the default route is sometimes suboptimal. For example, if the ProxySG satisfies a request and sends the client response traffic over the default route, the gateway router simply returns the traffic to the client router on the LAN side of the ProxySG. This causes unnecessary traffic between the switch and gateway, before the packet is finally received by the client router.




Figure 35–2 Effects of sending client response to the default gateway

To specify a direct route, a static route must be created, which requires the administrator to update the routing table every time a new route is needed.

The Return-to-Sender (RTS) option eliminates the need to create static routes by configuring the ProxySG to send response packets back to the same interface that received the request packet, entirely bypassing any routing lookup on the ProxySG. Essentially, the ProxySG stores the source Ethernet MAC address that the client's packet came from and sends all responses to that address.

The RTS interface mapping is updated each time a packet is received. For example, if there are two gateways and both of them send packets to the ProxySG, the packets are sent back to the last MAC address and interface that received the packet.

Note: Non-default static routes override RTS settings.

RTS can be configured in two ways, inbound or outbound. These two options can be enabled at the same time.

Inbound RTS affects connections initiated to the ProxySG by clients and is enabled by default in SGOS 5.4 and later. Inbound RTS configures the ProxySG to send SYN-ACK packets to the same interface that the SYN packet arrived on. All subsequent TCP/IP response packets are also sent to the same interface that received the request packet.

RTS inbound applies only to clients who are on a different subnet than the ProxySG. If clients are on the same subnet, interface routes are used.




Figure 35–3 Inbound RTS

Outbound RTS affects connections initiated by the ProxySG to origin servers. Outbound RTS causes the ProxySG to send ACK and subsequent packets to the same interface that the SYN-ACK packet arrived on. Outbound RTS requires a route to the client, either through a default gateway or static route.




Figure 35–4 Outbound RTS

Enabling Return-to-Sender

To enable RTS, use the **return-to-sender** command. For example:

```
#(config) return-to-sender inbound {disable | enable}
```

Enables or disables return-to-sender for inbound sessions.

```
#(config) return-to-sender outbound {disable | enable}
```

Enables or disables return-to-sender for outbound sessions.

DNS Verification

In transparent deployments, the ProxySG verifies the destination IP addresses provided by the client. This is known as *L2/L3 transparency*.

Note: The Trust Destination IP option overrides DNS verification. This option is recommended for acceleration deployments only. For more information about this option, see "[About Trusting the Destination IP Address Provided by the Client](#)" on page 133

For hostname-less protocols such as CIFS and FTP, the IP address can always be trusted. For other protocols, such as HTTP, RTSP, and MMS, which have a hostname that must be resolved, verification can be an issue. URL rewrites that modify the hostname also can cause verification to fail.

L2/L3 transparency is not supported in explicit proxy deployments, or if the destination IP addresses cannot be verified by the ProxySG. In these cases, you must configure static routes to hosts that are only accessible through gateways other than the default gateway.

Transparent ADN connections that are handed off to an application proxy (HTTP or MAPI, for example) can utilize L2/L3 transparency. Also, transparent ADN connections that are tunneled but not handed off can utilize the functionality.

Note: IM is not supported with trust client addressing. To support IM, proper routes must be configured for Internet access and IM client-to-client communication.

Chapter 36: Configuring Failover

Using IP address failover, you can create a redundant network for any explicit proxy configuration. If you require transparent proxy configuration, you can create software bridges to use failover. For information on creating software bridges, see ["Configuring a Software Bridge"](#) on page 1255.

Note: If you use the Pass-Through adapter for transparent proxy, you must create a software bridge rather than configuring failover. For information on using the Pass-Through adapter, see ["About the Pass-Through Adapter"](#) on page 1254.

Using a pool of IP addresses to provide redundancy and load balancing, Blue Coat moves these IP addresses among a group of machines.

Topics in this Section

This section includes information about the following topics:

- ❑ ["About Failover"](#)
- ❑ ["Configuring Failover Groups"](#) on page 808

About Failover

Failover allows a second machine to take over if a first machine (not just an interface card) fails, providing redundancy to the network through a master/slave relationship. In normal operations, the master (the machine whose IP address matches the group name) owns the address. The master sends keep alive messages (*advertisements*) to the slaves. If the slaves do not receive advertisements at the specified interval, the slave with the highest configured priority takes over for the master. When the master comes back online, the master takes over from the slave again.

The Blue Coat failover implementation resembles the Virtual Router Redundancy Protocol (VRRP) with the following exceptions:

- ❑ A configurable IP multicast address is the destination of the advertisements.
- ❑ The advertisement interval is included in protocol messages and is learned by the slaves.
- ❑ A virtual router identifier (VRID) is not used.
- ❑ Virtual MAC addresses are not used.
- ❑ MD5 is used for authentication at the application level.

Masters are elected, based on the following factors:

- ❑ If the failover mechanism is configured for a physical IP address, the machine owning the physical address have the highest priority. This is not configurable.
- ❑ If a machine is configured as a master using a virtual IP address, the master has a priority that is higher than the slaves.

When a slave takes over because the master fails, an event is logged in the event log. No e-mail notification is sent.

Configuring Failover Groups

Configuring failover groups is necessary to enable network redundancy on the ProxySG appliance.

Failover is enabled by completing the following tasks:

- ❑ Creating virtual IP addresses on each ProxySG appliance.
- ❑ Creating a failover group.
- ❑ Attach the failover group to the bridge configuration.
- ❑ Selecting a failover mode (parallel or serial - this can only be selected using the CLI).

Important: Configuring failover groups will not, in and of itself, enable failover in your ProxySG deployment. For additional information on configuring failover as well as conceptual information, see "[Configuring Failover](#)".


You also must decide which machine is the master and which machines are the slaves, and whether you want to configure explicit proxy or transparent proxy network.

When configuring the group, the master and all the systems in the group must have exactly the same failover configuration except for priority, which is used to determine the rank of the slave machines. If no priority is set, a default priority of 100 is used. If two appliances have equal priority, the one with the highest local IP address ranks higher.

Note: Configuring failover on an Application Data Network (ADN) is similar to configuring failover on other appliances, with the exception that you add a server subnet on multiple boxes instead of just one.

To configure failover:

1. Select the **Configuration > Network > Advanced > Failover** tab.
2. Click **New**. The Add Failover Group dialog displays.



3. Create a group using either a new IP address or an existing IP address. If the group has already been created, you cannot change the new IP address without deleting the group and starting over.

4. Configure group options:

- a. **Multicast address** refers to a Class D IP address that is used for multicast. It is not a virtual IP address.

Note: Class D IP addresses (224 to 239) are reserved for multicast. A Class D IP address has a first bit value of 1, second bit value of 1, third bit value of 1, and fourth bit value of 0. The other 28 bits identify the group of computers that receive the multicast message.

- b. **Relative Priority** refers to a range from 1-255 that is assigned to systems in the group. 255 is reserved for the system whose failover group ID equals the real IP address. (Optional) **Master** identifies the system with the highest priority (the priority value is greyed out).
- c. (Optional) **Advertisement Interval** refers to the length of time between advertisements sent by the group master. The default is 40 seconds. If the group master fails, the slave with the highest priority takes over (after approximately three times the interval value). The failover time of the group is controlled by setting this value.
- d. (Optional, but recommended) **Group Secret** refers to a password shared only with the group.


5. Select **enabled**.
6. Click **OK** to close the dialog.
7. Click **Apply**.

Viewing Failover Statistics

At any time, you can view statistics for any failover group you have configured on your system.

To view failover status:

1. Select **Statistics > System > Failover**.



The screenshot shows a window titled "Status" with a tabbed interface. The "Failover Group:" dropdown menu is set to "10.9.16.150". Below this, the "Failover status:" section displays the following information:

Multicast address:	224.1.2.3
Local address:	10.9.16.150
State:	MASTER
Flags:	R (Real IP)

2. From the drop-down list, select the group to view.

The information displayed includes the multicast address, the local address, the state, and any flags, where **V** indicates the group name is a virtual IP address, **R** indicates the group name is a physical IP address, and **M** indicates this machine can be configured to be the master if it is available.

Troubleshooting

An indication that there may be issues with the election of a master is if advertisements are not being sent or received by either of the systems in a failover group.

To troubleshoot, view statistics in the command line interface:

```
SGOS#(config) failover
SGOS#(config failover) view statistics
Failover Statistics
      Advertisements Received      : 0
      Advertisements Sent          : 0
      States Changes                : 0
      Bad Version                   : 0
      Bad Packet                    : 0
      Bad Checksum                  : 0
      Packet Too Short              : 0
      Bad Packet Header             : 0
      Invalid Group                 : 0
SGOS#(config failover)
```

If the statistics illustrate there may be a potential issue, debug further by running a PCAP on each ProxySG to verify the multicast packets are actually being sent. If not, verify the multicast address is configured correctly (**Configuration > Network > Advanced > Failover**). If both proxies are sending the multicast packets but not receiving them, it is possible that a switch/router is blocking multicast packets.

Chapter 37: Configuring DNS

This chapter describes various configuration tasks associated with Domain Name System (DNS) services. During first-time installation of the ProxySG appliance, you configured the IP address of a single primary DNS server. You can add one or more alternate DNS servers, as well as define custom DNS service groups.

Topics in this Chapter

This chapter includes the following topics:

- ❑ ["About DNS" on page 813](#)
- ❑ ["About Configuring DNS Server Groups" on page 815](#)
- ❑ ["Adding DNS Servers to the Primary or Alternate Group" on page 816](#)
- ❑ [""Promoting DNS Servers in a List"" on page 817](#)
- ❑ ["Creating a Custom DNS Group" on page 817](#)
- ❑ ["Deleting Domains" on page 818](#)
- ❑ ["Deleting DNS Groups and Servers" on page 818](#)
- ❑ ["Resolving Hostnames Using Name Imputing Suffixes" on page 819](#)
- ❑ ["Caching Negative Responses" on page 821](#)

About DNS

A hierarchical set of DNS servers comprises a Domain Name System. For each domain or sub-domain, one or more authoritative DNS servers publish information about that domain and the name servers of any domains that are under it.

Note: The DNS servers are configured in groups. For more information, see ["About Configuring DNS Server Groups" on page 815](#).

There are two types of queries, which are:

- ❑ Non-recursive, which means that a DNS server can provide a partial answer or return an error to the client
- ❑ Recursive, which means that the DNS server either fully answers the query or returns an error to the client

ProxySG Using Non-Recursive DNS

If you have defined more than one DNS server, the ProxySG uses the following logic to determine which servers are used to resolve a DNS host name and when to return an error to the client.

Note: Servers are always contacted in the order in which they appear in a group list.

- ❑ The ProxySG first checks all the DNS groups for a domain match, using domain-suffix matching to match a request to a group.
 - If there is a match, the servers in the matched group are queried until a response is received; no other DNS groups are queried.
 - If there is *no* match, the ProxySG selects the Primary DNS group.
- ❑ The ProxySG sends requests to DNS servers in the Primary DNS server group in the order in which they appear in the list. If a response is received from one of the servers in the Primary group, no attempts are made to contact any other Primary DNS servers.
- ❑ If none of the servers in the Primary group resolve the host name, the ProxySG sends requests to the servers in the Alternate DNS server group. (If no Alternate servers have been defined, an error is returned to the client.)
 - If a response is received from a server in the Alternate group list, there are no further queries to the Alternate group.
 - If a server in the Alternate DNS server group is unable to resolve the host name, an error is returned to the client, and no attempt is made to contact any other DNS servers.

Note: The Alternate DNS server is not used as a failover DNS server. It is only used when DNS resolution of the Primary DNS server returns a name error. If the query to each server in the Primary list times out, no alternate DNS server is contacted.

- If the ProxySG receives a referral (authoritative server information), DNS recursion takes over if it is enabled. See the next section, "[ProxySG Using Recursive DNS](#)" and "[When to Enable Recursive DNS](#)" on page 815.

Note: If the ProxySG receives a negative DNS response (a response with an error code set to `name error`), it caches that negative response. See "[Caching Negative Responses](#)" on page 821.

ProxySG Using Recursive DNS

If you have enabled recursive DNS, the ProxySG uses the following logic to determine how to resolve a DNS host name and when to return an error to the client.

- ❑ If the DNS server response does not contain an A record with an IP address but instead contains authoritative server information (a referral), the ProxySG follows all referrals until it receives an answer. If the ProxySG follows more than eight referrals, it assumes there is a recursion loop, aborts the request, and sends an error to the client.

When to Enable Recursive DNS

If you have a DNS server that cannot resolve all host names, it might return a list of authoritative DNS servers instead of a DNS A record that contains an IP address. To avoid this situation, configure the ProxySG to recursively query authoritative DNS servers.

To enable recursive DNS:

1. Select the **Configuration > Network > DNS > Groups** tab.
2. Select **Enable DNS Recursion**.
3. Click **Apply**.

To disable recursive DNS:

1. Select the **Configuration > Network > DNS > Groups** tab.
2. Clear **Enable DNS Recursion**.
3. Click **Apply**.

About Configuring DNS Server Groups

Customers with split DNS server configuration (for example, environments that maintain private internal DNS servers and external DNS servers) might choose to add servers to an Alternate DNS server group as well as to the Primary DNS server group. In addition, you can create custom DNS server groups.

In the ProxySG, internal DNS servers are placed in the Primary group, while external DNS servers (with the Internet information) populate the Alternate group.

The following rules apply to DNS server groups:

- ❑ You can add servers to the Primary and Alternate groups, but you cannot change the domain or add additional domains; these groups are defined at initial configuration.
- ❑ The Primary and Alternate DNS groups cannot be deleted.
- ❑ A custom DNS group must have at least one server in order to add domains.

About DNS Health Checks

Each time you add a DNS server to a group, the ProxySG automatically creates a DNS health check for that server IP address and uses a default configuration for the health check. For example, if you add a DNS server to a primary or alternate DNS group, the created health check has a default hostname of bluecoat.com. If you add a DNS server to a custom group, the longest domain name is used as the default hostname for the health check.

After you add DNS servers to a group, we recommend that you check the DNS server health check configurations and edit them as required. For complete details about configuring DNS server health checks, see "[About DNS Server Health Checks](#)" on page 1377."

Adding DNS Servers to the Primary or Alternate Group

This section discusses how to add DNS servers to a primary or alternate group. When you installed the ProxySG, you configured a Primary DNS server. If your deployment makes use of more than one DNS server, you can them to the Primary or Alternate server group; you can also delete DNS servers from the Primary group, but you cannot delete the group or change the domain or add additional domains—the group is defined at initial configuration.

If you are using the ProxySG in a mixed IPv4/IPv6 environment, you should configure both IPv4 and IPv6 DNS servers.

To add DNS servers to the Primary/Alternate group:

1. Select **Configuration > Network > DNS > Groups**.

Group Name	Servers	Domains
primary	10.2.2.100	*
alternate		*

☐ Enable DNS Recursion

2. Select a group (**primary** or **alternate**) and click **Edit**. The Edit DNS Forwarding Group dialog displays.

Edit DNS Forwarding Group

DNS Forwarding Group Settings

Group Name:

Servers	Domains
10.2.2.100	*
10.2.2.101	
10.2.2.102	

3. Enter the IPv4 or IPv6 address of each additional DNS server and click **OK**.
4. Click **Apply**.

See Also

- ❑ ["About DNS"](#)
- ❑ [""Promoting DNS Servers in a List""](#)

- ❑ "Creating a Custom DNS Group"
- ❑ "About Configuring DNS Server Groups"
- ❑ "Promoting DNS Servers in a List"

Creating a Custom DNS Group

Custom groups enable you to specify servers and domains for specific company needs (such as resolving internal or external hostnames) depending on how you have set up your primary and alternate DNS groups.

Valid DNS entry formats are:

```
example.com
www.example.com
```

Notes:

- ❑ You can create a maximum of 8 custom groups, and each custom group can contain a maximum of four DNS servers and eight domains.
- ❑ Groups do not accept wild cards, such as:

```
*.example.com
```

- ❑ Groups do *not* partially match domain names, such as:

```
*.example.com
.example.com
```

Further more:

```
exam.com
```

does not match queries for `www.example.com`.

- ❑ DNS record requirements have been relaxed, as discussed in RFC 2181. Review sections 10 and 11 for more information.

To create a custom group:

1. Select **Configuration > Network > DNS > Groups**. The list of DNS groups displays.
2. Click **New**. The Create DNS Forwarding Group dialog displays.

3. Enter a name for the DNS group.
4. Enter the servers (IPv4 or IPv6 addresses) and the domains for the group, and click **OK**. The custom group displays in the DNS Groups list.

5. Click **Save**.

See Also

- ❑ ["About DNS"](#)
- ❑ ["About Configuring DNS Server Groups"](#)
- ❑ ["Adding DNS Servers to the Primary or Alternate Group"](#)
- ❑ [""Promoting DNS Servers in a List""](#)
- ❑ ["Promoting DNS Servers in a List"](#)

Deleting Domains

If a domain becomes defunct, you can easily delete it from a DNS group. In addition, you need to delete all domains associated with the last server in any DNS group before you can delete the server.

To delete domains:

1. Select **Configuration > Network > DNS > Groups**. The list of DNS groups displays.
2. Select the DNS group in the list and click **Edit**. The Edit DNS Forwarding Group dialog displays.
3. Delete domains, and click **OK**.
4. Click **Apply**.

See Also

["Deleting DNS Groups and Servers"](#)

Deleting DNS Groups and Servers

The following list describes the specific rules that apply when deleting DNS groups and servers.

- ❑ You cannot delete the Primary or Alternate DNS group; you can only delete a custom DNS group.
- ❑ You cannot delete the last server in any DNS group while there are still domains that reference that group; doing so returns an error message.

To delete a DNS server:

1. Select **Configuration > Network > DNS > Groups**.
2. Select the DNS group from which to delete a server, and click **Edit**. The **Edit DNS Forwarding Group** dialog displays.
3. Delete the server, then click **OK**.
4. Click **Apply**.

To delete a custom DNS group:

1. Select **Configuration > Network > DNS > Groups**.

2. Select the custom DNS group to delete, and click **Delete**. A dialog box displays, asking you to confirm your choice.
3. Click **OK** to delete the group.

See Also

- ❑ ["Deleting Domains"](#)
- ❑ ["Promoting DNS Servers in a List"](#)

Promoting DNS Servers in a List

Using the CLI, you can promote DNS servers in the list for any DNS forwarding group.

To promote DNS servers in a list:

```
#(config dns forwarding) edit group_alias
```

This changes the prompt to:

```
#(config dns forwarding group)
#(config dns forwarding group) promote server_ip #
```

This promotes the specified server IP address in the DNS server list the number of places indicated. You must use a positive number. If the number is greater than the number of servers in the list, the server is promoted to the first entry in the list.

See Also

- ❑ ["Adding DNS Servers to the Primary or Alternate Group"](#)
- ❑ [""Promoting DNS Servers in a List""](#)
- ❑ ["Creating a Custom DNS Group"](#)
- ❑ ["Deleting DNS Groups and Servers"](#)

Resolving Hostnames Using Name Imputing Suffixes

The ProxySG queries the original hostname before checking imputing suffixes *unless* there is no period in the hostname. If there is no period in the hostname, imputing is applied first.

The ProxySG uses name imputing to resolve hostnames based on a partial name specification (DNS name imputing suffix). When the ProxySG submits a hostname to the DNS server, the DNS server resolves the hostname to an IP address.

The ProxySG then tries each entry in the name-imputing suffixes list until the name is resolved or it reaches the end of the list. If by the end of the list the name is not resolved, the ProxySG returns a DNS failure.

For example, if the name-imputing list contains the entries `example.com` and `com`, and a user submits the URL `http://www.eeddept`, the ProxySG resolves the host names in the following order.


www.eeddept
www.eeddept.example.com
www.eeddept.com

Adding and Editing DNS Name Imputing Suffixes

Using name imputing suffixes is particularly useful for a company's internal domains. For example, it enables you to simply enter `webServer` rather than the more elaborate `webServer.inOurInternalDomain.ForOurCompany.com`. Also, this resolves any problem with external `root` servers being unable to resolve names that are internal only. The ProxySG supports up to 30 name imputing suffixes.

To add names to the imputing list:


1. Select the **Configuration > Network > DNS > Imputing** tab.
2. Click **New**. The Add List Item dialog displays.



3. Enter the DNS name imputing suffix and click **OK**.
The name displays in the DNS name imputing suffixes list.
4. Click **Apply**.

To edit DNS name imputing suffixes:

1. Select the **Configuration > Network > DNS > Imputing** tab.
2. Select a name in the list and click **Edit**. The Edit List Item dialog displays.



3. Edit the name imputing suffix as required and click **OK**.
4. Click **Apply**.

Changing the Order of DNS Name Imputing Suffixes

The ProxySG uses imputing suffixes according to the list order. You can organize the list of suffixes so the preferred suffix displays at the top of the list.

Note: This functionality is only available through the Management Console. You cannot configure it using the CLI.

To change the order of DNS name imputing suffixes:

1. Select **Configuration > Network > DNS > Imputing**.
2. Select the imputing suffix to promote or demote.
3. Click **Promote entry** or **Demote entry**, as appropriate.
4. Click **Apply**.

Caching Negative Responses

By default, the ProxySG caches negative DNS responses sent by a DNS server. You can configure the ProxySG to set the time-to-live (TTL) value for a negative DNS response to be cached. You can also disable negative DNS response caching.

Note: The ProxySG generates more DNS requests when negative caching is disabled.

The ProxySG supports caching of both type A and type PTR DNS negative responses.

This functionality is only available through the CLI. You cannot configure DNS negative caching through the Management Console.

To configure negative caching TTL values:

From the `(config)` prompt:

```
SGOS#(config) dns negative-cache-ttl-override seconds
```

where *seconds* is any integer between 0 and 600.

Setting the TTL value to 0 seconds disables negative DNS caching; setting the TTL setting to a non-zero value overrides the TTL value from the DNS response.

To restore negative caching defaults:

From the `(config)` prompt:

```
SGOS#(config) dns no negative-cache-ttl-override
```


Chapter 38: Virtual IP Addresses

This chapter discusses the uses of Virtual IP (VIP) addresses and how to create them.

Virtual IP addresses are addresses assigned to a system (but not an interface) that are recognized by other systems on the network.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ ["Uses of a VIP"](#) on page 823
- ❑ ["Creating a VIP"](#) on page 823
- ❑ ["Deleting a VIP"](#) on page 824

Uses of a VIP

VIP addresses have several uses:

- ❑ Assign multiple identities to a system on the same or different network, partitioning the box in to separate logical entities for resource sharing or load sharing.
- ❑ Create an HTTPS Console to allow multiple, simultaneous, secure connections to the system.
- ❑ Direct authentication challenges to different realms.
- ❑ Set up failover among multiple ProxySG appliances on the same subnet.

Note: For information on creating an HTTPS Console, see ["Managing the HTTPS Console \(Secure Console\)"](#) on page 1270; for information on using VIPs with authentication realms, see ["About Origin-Style Redirection"](#) on page 910; to use VIPs with failover, see ["Configuring Failover"](#) on page 807.

Creating a VIP

You can create up to 255 VIPs through the Management Console. To create more VIPs, use the `#(config)virtual-ip address <IP_address>` command.

To create a VIP:

1. In the Management Console, select **Configuration > Network > Advanced > VIPs**.
2. Click **New**.
3. Enter the virtual IP address you want to use. It can be any IP address, except a multicast address. (A multicast address is a group address, not an individual IP address.)

Note: You cannot create a VIP address that is the IP address used by the origin content server. You must assign a different address on the ProxySG, and use DNS or forwarding to point to the origin content server's real IP address.

4. Click **OK**.

5. Click **Apply**.

The VIP address can now be used.

Deleting a VIP

To delete a VIP:

1. In the Management Console, select **Configuration > Network > Advanced > VIPs**.
2. Select the VIP you want to delete.
3. Click **Delete**.
4. Click **OK** on the Confirm Delete dialog that appears.

Chapter 39: Configuring Private Networks

This chapter describes how the ProxySG interacts in internal, or private, networks.

Topics in this Chapter

This chapter includes information on the following topics:

- ❑ ["About Private Networks" on page 825](#)
- ❑ ["Default Private Subnets on the ProxySG" on page 826](#)
- ❑ ["Configuring Private Subnets" on page 826](#)
- ❑ ["Configuring Private Domains" on page 827](#)
- ❑ ["Using Policy On Configured Private Networks" on page 828](#)

About Private Networks

A private network is an internal network that uses private IP addresses, which are usually not routed over the public Internet. For example, your intranet that forms an important component of internal communication and collaboration, could have private websites — private domains and private subnets.

This security feature allows you to control private information within your network. Any private host that is configured on the ProxySG is identified as internal traffic and *dynamic categorization* or WebPulse is not performed on that host.

Further, if you configure a private domain that includes hosts with routable IP addresses on the ProxySG, you can use policy to suppress information. For example, you can suppress sensitive information like the HTTP `Referer` information from being sent over the internet.

Also, if you have a DMZ network that includes hosts with routable IP addresses that might be accessed from the Internet, you can configure these IP addresses as a part of your private network. You can then create policy to restrict access to your private network. Thereby, configuring a private network allows you to enhance performance and security within your network.

Default Private Subnets on the ProxySG

The ProxySG is pre-configured with private subnets that use non-routable IP addresses. These non-routable addresses provide additional security to your private network because packets using IP addresses within this range are rejected by Internet routers.

Table 39–1 Private Subnets on the ProxySG

Pre-configured Private Subnets	Details
0.0.0.0/8	Source Hosts on This Network
10.0.0.0/8	Private Networks Class A
127.0.0.0/8	Internet Host Loopback Address
169.254.0.0/16	"Link Local" Block
172.16.0.0/12	Private Networks Class B
192.168.0.0/16	Private Networks Class C
224.0.0.0/3	Multicast + Reserved


The ProxySG allows you to delete subnets from this list or add private subnets to this list, see ["Configuring Private Subnets"](#) on page 826 to configure private subnets. To configure private domains, see ["Configuring Private Domains"](#) on page 827.

Configuring Private Subnets

A private subnet consists of IP addresses that are generally not directly accessible from the Internet.

To Add a Private Subnet on the ProxySG:

1. Select the **Configuration > Network > Private Network > Private Subnets** tab.



2. Click **Add**. The **Add Private Subnet** dialog displays.
3. Enter the **IP Address** or the **Subnet Prefix**, and the **Subnet Mask** of the private subnet.
4. Click **OK**.

To Remove a Private Subnet on the ProxySG:

1. On the **Configuration > Network > Private Networks** tab, select the private subnet to delete.
2. Click **Remove**.

To Restore the Default Private Subnets Configured on the ProxySG:

On the **Configuration > Network > Private Networks** tab, click **Set to Default**. The ProxySG reverts to the default list of non-routable IP addresses.

See Also

["Configuring Private Networks"](#)

["Default Private Subnets on the ProxySG"](#)

["Using Policy On Configured Private Networks"](#)


Configuring Private Domains

A domain name is an easy to remember name for an IP address. For example, if the private IP address 10.0.0.2 has the hostname `intranet.xyz.com`, you can define the domain `xyz.com` within the **Private Domain** list for your network.

If you then implement policy that restricts access logging or transferring of sensitive information, the interaction between a client and any host on the domain `xyz.com` can be kept private, that is any information pertaining to this private network is not sent over the public Internet. For details on implementing policy for the configured private network, see ["Using Policy On Configured Private Networks"](#) on page 828.

To Add a Private Domain on the ProxySG:

1. Select the **Configuration > Network > Private Network > Private Domains** tab.



2. Click **Add**. The **Add Private Domains** dialog displays.
3. Enter the internal domain information. Add one domain per line.
4. Click **Add**.
5. Click **Close**.

To Delete One or More Private Domain(s) on the ProxySG

1. On the **Configuration > Network > Private Networks > Private Domains** tab, select the private domain to delete.
2. Click **Remove**, to delete the selected domain.
3. Or, click **Clear All** to delete all private domains configured on the ProxySG.

See Also

- ❑ ["Configuring Private Networks"](#)
- ❑ ["Default Private Subnets on the ProxySG"](#)
- ❑ ["Configuring Private Subnets"](#)
- ❑ ["Using Policy On Configured Private Networks"](#)

Using Policy On Configured Private Networks

Blue Coat Policy allows administrators to create and apply flexible policies in their network. This section includes information on the Content Policy Language (CPL) gestures that are available for testing private hosts and on how these gestures can be used to create policy. The following topics are covered in this section:

- ❑ ["CPL Gestures for Validating Private Hosts" on page 828](#)
- ❑ ["Restricting Access Logging for Private Subnets" on page 829](#)
- ❑ ["Stripping Referer Header for Internal Servers" on page 829](#)

CPL Gestures for Validating Private Hosts

The following Content Policy Language gestures are available for testing private hosts that are configured on the ProxySG:

- `url.host.is_private` compares whether the host name in a request URL belongs to a private domain configured on the ProxySG.
- `request.header.referer.url.host.is_private` examines whether the Referer header in an HTTP request belongs to a private domain configured on the ProxySG.
- `server_url.host.is_private` compares whether the host name in a server URL belongs to a private domain configured on the ProxySG. The server URL is the URL in the request issued by the ProxySG to the OCS. The server URL is usually the same as the request URL, but it can be different if URL rewriting is implemented on the ProxySG.

You can use these gestures to create policy and to manage exceptions. The following example creates a whitelist for virus scanning and demonstrates the use of the `url.host.is_private` gesture.

```
define condition extension_low_risk
url.extension=(asf,asx,gif,jpeg,mov,mp3,ram,rm,smi,smil,swf,txt,wax,wma,wmv,wvx)
end
```

```
<cache>
    condition=extension_low_risk response.icap_service
    (icap_server, fail_open)
    response.icap_service(icap_server, fail_closed)
; exception
<cache>
    url.host.is_private=yes response.icap_service(no)
```

The task flow for creating this policy is:

- a. Create a list with the `define condition` gesture. Definitions allow you to bind a set of conditions or actions to your list.
- b. For the list defined, assign the file types that you regard as low risk for viruses.
- c. Create a condition for web content, in the `<cache>` layer, that specifies the ICAP response service to fail open for the low-risk file types as defined in the `extension_low_risk` list, while all other files will fail closed until the scan is completed.
- d. Create an exception that in the `<cache>` layer that exempts scanning of all responses from internal hosts, since internal hosts are considered secure.

For more information on using policy and for details on CPL gestures, refer to the *Content Policy Language Guide*.

Restricting Access Logging for Private Subnets

Since a private subnet belongs to an internal network, you might decide not to log requests made to private servers in the access log. The following policy example enables you to log access to public sites only.

```
<Proxy>
url.host.is_private=yes access_log(no)
```

Stripping Referer Header for Internal Servers

If a server in your private network refers or links to a public website, you can remove or suppress sensitive information like the HTTP `Referer` details. Stripping the header allows you to withhold information about the web servers in your private network. To strip the `Referer` header, use the following policy:

```
<Proxy>
request.header.Referer.url.host.is_private=yes action.HideReferer(yes)

define action HideReferer
delete(request.header.Referer)
end action HideReferer
```


Chapter 40: Managing Routing Information Protocols (RIP)

This chapter discusses the Routing Information Protocol (RIP), which is designed to select the fastest route to a destination. RIP support is built into the ProxySG appliance, and is configured by creating and installing an RIP configuration text file onto the device.

The Blue Coat RIP implementation also supports advertising default gateways. Default routes added by RIP are treated the same as the static default routes; that is, the default route load balancing schemes apply to the default routes from RIP as well.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ "Installing RIP Configuration Files" on page 831
- ❑ "Configuring Advertising Default Routes" on page 832
- ❑ "RIP Commands" on page 833
- ❑ "RIP Parameters" on page 834
- ❑ "ProxySG-Specific RIP Parameters" on page 836
- ❑ "Using Passwords with RIP" on page 837

Installing RIP Configuration Files

No RIP configuration file is shipped with the appliance. For commands that can be entered into the RIP configuration file, see "RIP Commands" on page 833.

After creating an RIP configuration file, install it using one of the following methods:

- ❑ Using the Text Editor, which allows you to enter settings (or copy and paste the contents of an already-created file) directly onto the appliance.
- ❑ Creating a local file on your local system; the ProxySG can browse to the file and install it.
- ❑ Using a remote URL, where you place an already-created file on an FTP or HTTP server to be downloaded to the ProxySG.
- ❑ Using the CLI `inline rip-settings` command, which allows you to paste the RIP settings into the CLI.
- ❑ Using the CLI `rip` commands, which require that you place an already-created file on an FTP or HTTP server and enter the URL into the CLI. You can also enable or disable RIP with these commands.

To install an RIP configuration file:

Note: When entering RIP settings that affect current settings (for example, when switching from ripv1 to ripv2), disable RIP before you change the settings; re-enable RIP when you have finished.

1. Select **Configuration > Network > Routing > RIP**.
2. To display the current RIP settings, routes, or source, click one or all of the **View RIP** buttons.
3. In the **Install RIP Setting from** drop-down list, select the method used to install the routing table; click **Install**.
 - **Remote URL:**

Enter the fully-qualified URL, including the filename, where the routing table is located. To view the file before installing it, click **View**. Click **Install**. To view the installation results, click **Results**; close the window when you are finished. Click **OK**.
 - **Local File:**

Click **Browse** to display the Local File Browse window. Browse for the file on the local system. Open it and click **Install**. When the installation is complete, a results window opens. View the results and close the window.
 - **Text Editor:**

The current configuration is displayed in installable list format. You can customize it or delete it and create your own. Click **Install**. When the installation is complete, a results window opens. View the results, close the window, and click **OK**.
4. Click **Apply**.
5. Select **Enable RIP**.
6. Click **Apply**.

Configuring Advertising Default Routes

Default routes advertisements are treated the same as the static default routes; that is, the default route load balancing schemes also apply to the default routes from RIP.

By default, RIP ignores the default routes advertisement. You can change the default from disable to enable and set the preference group and weight through the CLI only.

To enable and configure advertising default gateway routes:

1. At the (config) command prompt:

```
SGOS#(config) rip default-route enable
SGOS#(config) rip default-route group group_number
SGOS#(config) rip default-route weight weight_number
```


Where *group_number* defaults to 1, and *weight_number* defaults to 100, the same as the static default route set by the `ip-default-gateway` command.

2. (Optional) To view the default advertising routes, enter:

```
SGOS#(config) show rip default-route
RIP default route settings:
Enabled:                               Yes
Preference group:                      3
Weight:                               30
```

RIP Commands

You can place any of the commands below into a Routing Information Protocol (RIP) configuration text file. You cannot edit a RIP file through the command line, but you can overwrite a RIP file using the `inline rip-settings` command.

After the file is complete, place it on an HTTP or FTP server accessible to the ProxySG and download it.

Note: RIP parameters are accepted in the order that they are entered. If a RIP parameter is added, it is appended to the default RIP parameters. If a subsequent parameter conflicts with a previous parameter, the most recent one is used.

net

```
net Nname[/mask] gateway Gname metric Value {passive | active |
external}
```

Table 40–1 net Commands

Parameters	Description
<i>Nname</i>	Name of the destination network. It can be a symbolic network name, or an Internet address specified in dot notation.
<i>/mask</i>	Optional number between 1 and 32 indicating the netmask associated with <i>Nname</i> .
<i>Gname</i>	Name or address of the gateway to which RIP responses should be forwarded.
<i>Value</i>	The hop count to the destination host or network. A net <i>Nname/32</i> specification is equivalent to the host <i>Hname</i> command.
passive active external	Specifies whether the gateway is treated as passive or active, or whether the gateway is external to the scope of the RIP protocol.

host

```
host Hname gateway Gname metric Value {passive | active | external}
```

Table 40–2 host Commands

Parameters	Description
<i>Hname</i>	Name of the destination network. It can be a symbolic network name, or an Internet address specified in dot notation.
<i>Gname</i>	Name or address of the gateway to which RIP responses should be forwarded. It can be a symbolic network name, or an Internet address specified in dot notation.
<i>Value</i>	The hop count to the destination host or network. A net <i>Nname</i> /32 specification is equivalent to the host <i>Hname</i> command.
passive active external	Specifies whether the gateway is treated as passive or active, or whether the gateway is external to the scope of the RIP protocol.

RIP Parameters

Lines that do not start with net or host commands *must* consist of one or more of the following parameter settings, separated by commas or blank spaces:

Table 40–3 RIP Parameters

Parameters	Description
if=[0 1 2 3]	Specifies that the other parameters on the line apply to the interface numbered 0,1,2, or 3 in SGOS terms.
passwd=XXX	Specifies an RIPv2 password included on all RIPv2 responses sent and checked on all RIPv2 responses received. The password must not contain any blanks, tab characters, commas or '#' characters.
no_ag	Turns off aggregation of subnets in RIPv1 and RIPv2 responses.
no_super_ag	Turns off aggregation of networks into supernets in RIPv2 responses.
passive	Marks the interface to not be advertised in updates sent through other interfaces, and turns off all RIP and router discovery through the interface.
no_rip	Disables all RIP processing on the specified interface.
no_ripv1_in	Causes RIPv1 received responses to be ignored.
no_ripv2_in	Causes RIPv2 received responses to be ignored.
ripv2_out	Turns off RIPv1 output and causes RIPv2 advertisements to be multicast when possible.
ripv2	Is equivalent to no_ripv1_in and no_ripv1_out. This parameter is set by default.
no_rdisc	Disables the Internet Router Discovery Protocol. This parameter is set by default.

Table 40–3 RIP Parameters (Continued)

Parameters	Description
<code>no_solicit</code>	Disables the transmission of Router Discovery Solicitations.
<code>send_solicit</code>	Specifies that Router Discovery solicitations should be sent, even on point-to-point links, which by default only listen to Router Discovery messages.
<code>no_rdisc_adv</code>	Disables the transmission of Router Discovery Advertisements.
<code>rdisc_adv</code>	Specifies that Router Discovery Advertisements should be sent, even on point-to-point links, which by default only listen to Router Discovery messages.
<code>bcast_rdisc</code>	Specifies that Router Discovery packets should be broadcast instead of multicast.
<code>rdisc_pref=N</code>	Sets the preference in Router Discovery Advertisements to the integer N.
<code>rdisc_interval=N</code>	Sets the nominal interval with which Router Discovery Advertisements are transmitted to N seconds and their lifetime to 3*N.
<code>trust_gateway=rname</code>	Causes RIP packets from that router and other routers named in other <code>trust_gateway</code> keywords to be accept, and packets from other routers to be ignored.
<code>redirect_ok</code>	Causes RIP to allow ICMP Redirect messages when the system is acting as a router and forwarding packets. Otherwise, ICMP Redirect messages are overridden.

ProxySG-Specific RIP Parameters

The following RIP parameters are unique to ProxySG configurations:

Table 40–4 ProxySG-Specific RIP Parameters

Parameters	Description
supply_routing_info -or- advertise_routes	-s option: Supplying this option forces routers to supply routing information whether it is acting as an Internetwork router or not. This is the default if multiple network interfaces are present or if a point-to-point link is in use. -g option: This flag is used on Internetwork routers to offer a route to the 'default' destination. This is typically used on a gateway to the Internet, or on a gateway that uses another routing protocol whose routes are not reported to other local routers. -h option: Suppress_extra_host_routes advertise_host_route -m option: Advertise_host_route on multi-homed hosts -A option: Ignore_authentication //
no_supply_routing_info	-q option: opposite of -s.
no_rip_out	Disables the transmission of all RIP packets. This setting is the default.
no_ripv1_out	Disables the transmission of RIPv1 packets.
no_ripv2_out	Disables the transmission of RIPv2 packets.
rip_out	Enables the transmission of RIPv1 packets.
ripv1_out	Enables the transmission of RIPv1 packets.
rdisc	Enables the transmission of Router Discovery Advertisements.
ripv1	Causes RIPv1 packets to be sent.
ripv1_in	Causes RIPv1 received responses to be handled.

Using Passwords with RIP

The first password specified for an interface is used for output. All passwords pertaining to an interface are accepted on input. For example, with the following settings:

```
if=0 passwd=aaa  
if=1 passwd=bbb  
passwd=ccc
```

Interface 0 accepts passwords `aaa` and `ccc`, and transmits using password `aaa`. Interface 1 accepts passwords `bbb` and `ccc`, and transmits using password `bbb`. The other interfaces accept and transmit the password `ccc`.

Chapter 41: SOCKS Gateway Configuration

This chapter discusses the Blue Coat implementation of SOCKS, which includes the following:

- ❑ A SOCKS proxy server that supports both SOCKSv4/4a and SOCKSv5, running on the ProxySG appliance.
- ❑ Support for forwarding through SOCKS gateways.

To configure a SOCKS proxy server on the ProxySG, see [Chapter 14: "Managing a SOCKS Proxy"](#) on page 305. To use SOCKS gateways when forwarding traffic, continue with this chapter.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ [Section A: "Configuring a SOCKS Gateway"](#) on page 840.
- ❑ [Section B: "Using SOCKS Gateways Directives with Installable Lists"](#) on page 849.

Section A: Configuring a SOCKS Gateway

The following topics in this section discuss how to configure a SOCKS gateway, groups, defaults, and the default sequence:

- ❑ ["About SOCKS Gateways"](#)
- ❑ ["Adding a SOCKS Gateway" on page 840](#)
- ❑ ["Creating SOCKS Gateway Groups" on page 843](#)
- ❑ ["Configuring Global SOCKS Defaults" on page 845](#)
- ❑ ["Configuring the SOCKS Gateway Default Sequence" on page 847](#)

About SOCKS Gateways

SOCKS servers provide application-level firewall protection for an enterprise.

SOCKS gateways (forwarding) can use installable lists for configuration.

Configure the installable list using directives. You can also use the Management Console or the CLI to create a SOCKS gateways configuration. Using the Management Console is the easiest method.


See Also

- ❑ ["Adding a SOCKS Gateway" on page 840](#)
- ❑ ["Creating SOCKS Gateway Groups" on page 843](#)
- ❑ ["Configuring Global SOCKS Defaults" on page 845](#)
- ❑ ["Configuring the SOCKS Gateway Default Sequence" on page 847](#)

Adding a SOCKS Gateway

To configure a SOCKS gateway:

1. Select the **Configuration > Forwarding > SOCKS Gateways > SOCKS Gateways** tab.
2. Click **New** to create a new SOCKS gateway.



3. Configure the SOCKS gateway as follows:
 - a. **Alias:** Give the gateway a meaningful name.

Note: SOCKS gateway aliases cannot be CPL keywords, such as `no`, `default`, `forward`, or `socks_gateways`.

- b. **Host:** Add the IP address or the host name of the gateway where traffic is directed. The host name must DNS resolve.
- c. **Port:** The default is 1080.
- d. **SOCKS version:** Select the version that the SOCKS gateway can support from the drop-down list. Version 5 is recommended.
- e. **Username** (Optional, and only if you use version 5) The username of the user on the SOCKS gateway. The username already must exist on the gateway. If you have a username, you must also set the password.
- f. **Set Password:** The plaintext password or encrypted password of the user on the SOCKS gateway. The password must match the gateway's information. The password can be up to 64 bytes long. Passwords that include spaces must be within quotes.

You can enter an encrypted password (up to 64 bytes long) either through the CLI or through installable list directives.

- g. In the **Load Balancing and Host Affinity** section, select the load balancing method from the drop-down list. **Global default** (configured on the **Configuration > Forwarding > Global Defaults** tab), sets the default for all SOCKS gateways on the system. You can also specify the load balancing method for this system: **Least Connections** or **Round Robin**, or you can disable load balancing by selecting **None**.
- h. In the **Host affinity methods** drop-down list, select the method you want to use:
 - **HTTP:** The default is to use the **Global Defaults**. Other choices are **None**, which disables host affinity, **Accelerator Cookie**, which places a cookie in the response to the client, and **Client IP Address**, which uses the client IP address to determine which upstream SOCKS gateway was last used.

By default, SOCKS treats all incoming requests destined to port 80 as HTTP, allowing the usual HTTP policy to be performed on them, including ICAP scanning. If the SOCKS connection is being made to a server on another port, write policy on the ProxySG to match on the server host and port and specify that it is HTTP using SOCKS.

- **SSL:** The default is to use the **Global Defaults**. Other choices are **None**, which disables host affinity, **Accelerator Cookie**, which places a cookie in the response to the client, and **Client IP Address**, which uses the client IP address to determine which group member was last used. In addition, you can select **SSL Session ID**, used in place of a cookie or IP address, which extracts the SSL session ID name from the connection information.

- **Other:** Applies to any traffic that is not HTTP, terminated HTTPS, or intercepted HTTPS. You can attempt load balancing of any of the supported traffic types in forwarding and this host affinity setting can be applied as well. For example, you could load balance a set of TCP tunnels and apply the **Other** host affinity (client IP only).

The default is to use **Global Defaults**. Other choices are **None**, which disables host affinity, and **Client IP Address**, which uses the client IP address to determine which group member was last used.

- Click **OK** to close the dialog.
- Click **Apply**.


Creating SOCKS Gateway Groups

To create groups:

An existing gateway can belong to none, one, or more groups as desired (it can only belong once to a single group, however).

1. Select the **Configuration > Forwarding > SOCKS Gateways > SOCKS Gateway Groups** tab.

- Click **New**. The Add SOCKS Gateway Group dialog displays.



- To create an alias group, highlight the hosts and groups you want grouped, and click **Add**.
- Give the new group a meaningful name.
- In the **Load Balancing and Host Affinity** section, select the load balancing method from the drop-down list. **Global default** (configured on the **Configuration > Forwarding > SOCKS Gateways > Global Defaults** tab), sets the default for all forwarding hosts on the system. You can also specify the load balancing method for this system: **Least Connections**, **Round Robin**, **Domain Hash**, **URL Hash**, or you can disable load balancing by selecting **None**.

6. In the **Host affinity methods** drop-down lists, select the method you want to use. Refer to the previous procedure for details on methods. You are selecting between the resolved IP addresses of all of the hosts in the group, not the resolved IP addresses of an individual host.
 - **HTTP**: The default is to use the **Global Defaults**. Other choices are **None**, which disables host affinity, **Accelerator Cookie**, which places a cookie in the response to the client, and **Client IP Address**, which uses the client IP address to determine which group member was last used.
 - **SSL**: The default is to use the **Global Defaults**. Other choices are **None**, which disables host affinity, **Accelerator Cookie**, which places a cookie in the response to the client, and **Client IP Address**, which uses the client IP address to determine which group member was last used. In addition, you can select **SSL Session ID**, used in place of a cookie or IP address, which extracts the SSL session ID name from the connection information.
 - **Other**. Applies to any traffic that is not HTTP, terminated HTTPS, or intercepted HTTPS. You can attempt load balancing of any of the supported traffic types in forwarding and this host affinity setting can be applied as well. For example, you could load balance a set of TCP tunnels and apply the **Other** host affinity (client IP only).

The default is to use **Global Defaults**. Other choices are **None**, which disables host affinity, and **Client IP Address**, which uses the client IP address to determine which group member was last used.
7. Click **OK** to close the dialog.
8. Click **Apply**.

Configuring Global SOCKS Defaults

The global defaults apply to all SOCKS gateways hosts and groups unless the settings are specifically overwritten during host or group configuration.

To configure global defaults:

1. Select the **Configuration > Forwarding > SOCKS Gateways > Global Defaults** tab.

SOCKS Gateways	SOCKS Gateway Groups	Global Defaults	Default Sequence	Install SOCKS
General Settings				
If no forwarding host is available: ☐ Connect directly (fail open) ☒ Deny the request (fail closed)				
Load Balancing and Host Affinity				
Load balancing methods: Forwarding hosts: Round Robin Forwarding groups: Round Robin Host affinity methods: HTTP: None SSL: None Other: None Host affinity timeout: 30 minutes				

2. Determine how you want connections to behave if the health checks fail:
Connect Directly (fail open) or **Deny the request (fail closed)**. Note that failing open is an insecure option. The default is to fail closed. This option can be overridden by policy, if it exists.
3. In the **Global Load Balancing and Host Affinity** area:
 - a. Configure **Load Balancing methods**:
 - **SOCKS hosts**: Specify the load balancing method for all forwarding hosts unless their configuration specifically overwrites the global settings. You can choose **Least Connections** or **Round Robin**, or you can disable load balancing by selecting **None**. **Round Robin** is specified by default.
 - **SOCKS groups**: Specify the load balancing method for all forwarding groups unless their configuration specifically overwrites the global settings. You can choose to hash the domain or the full URL. You can also choose **Least Connections**, **Round Robin**, **Domain Hash**, **URL Hash**, and you can disable load balancing by selecting **None**. **Round Robin** is specified by default.

b. Configure **Global Host Affinity** methods:

- **HTTP:** The default is to use **None**, which disables host affinity. Other choices are **Accelerator Cookie**, which places a cookie in the response to the client, and **Client IP Address**, which uses the client IP address to determine which group member was last used.
- **SSL:** The default is to use **None**, which disables host affinity. Other choices are **Accelerator Cookie**, which places a cookie in the response to the client, and **Client IP Address**, which uses the client IP address to determine which group member was last used, and **SSL Session ID**, used in place of a cookie or IP address, which extracts the SSL session ID name from the connection information.
- **Other:** **Other** applies to any traffic that is not HTTP, terminated HTTPS, or intercepted HTTPS. You can attempt load balancing of any of the supported traffic types in forwarding and this host affinity setting can be applied as well. For example, you could load balance a set of TCP tunnels and apply the **Other** host affinity (client IP only).

The default is to use **None**, which disables host affinity. You can also choose **Client IP Address**, which uses the client IP address to determine which group member was last used.

- c. **Host Affinity Timeout:** This is the amount of time a user's IP address, SSL ID, or cookie remains valid. The default is 30 minutes, meaning that the IP address, SSL ID or cookie must be used once every 30 minutes to restart the timeout period.

4. Click **Apply**.

Configuring the SOCKS Gateway Default Sequence

The default sequence defines what SOCKS gateways to use when no policy is present to specify something different. The system uses the first host or group in the sequence that is healthy, just as it does when a sequence is specified through policy. Only one default sequence is allowed. All members must be pre-existing hosts, and no member can be in the group more than once.

A default failover sequence allow healthy hosts to take over for an unhealthy host (one that is failing its DNS Resolution or its health check). The sequence specifies the order of failover, with the second host taking over for the first host, the third taking over for the second, and so on.

If all hosts are unhealthy, the operation fails either open or closed, depending upon your settings.

This configuration is usually created and managed through policy. If no SOCKS-gateways policy applies, you can create a default sequence using policy. This single default sequence consists of a single default host (or group) plus one or more hosts to use if the preceding ones are unhealthy.

To create the default sequence:

Note: Traffic is forwarded to the first member of the list until it fails, then traffic is sent to the second member of list until it fails or the first member becomes healthy again, and so on.

1. Select the **Configuration > Forwarding > SOCKS Gateways > Default Sequence** tab.

The screenshot shows the 'Default Sequence' configuration window. It features a tabbed interface with 'SOCKS Gateways', 'SOCKS Gateway Groups', 'Global Defaults', and 'Default Sequence'. The 'Default Sequence' tab is selected, displaying a 'Default Fail-Over Sequence' section. This section is divided into two main areas: 'Available Aliases' and 'Selected Aliases'. The 'Available Aliases' area contains a list with one entry, 'test19'. The 'Selected Aliases' area is currently empty. Between these two areas are four buttons: 'Add >>', 'Promote', 'Demote', and '<< Remove'.

2. The available aliases (host and group) display in the **Available Aliases** pane. To select an alias, highlight it and click **Add**.

Note: Any host or group in the default sequence is considered in use by policy. As a result, if you try to delete a host or group while it is in the default sequence, you receive an error message. You must remove the host/group from the sequence first, then delete the host or group.

3. You can use the **Promote** and **Demote** buttons to change the order of the hosts and groups in the sequence after you add them to the **Selected Aliases** pane.
4. Click **Apply**.

Statistics

SOCKS gateways statistics are available through the **Statistics > Advanced > SOCKS Gateways** menu item.

Section B: Using SOCKS Gateways Directives with Installable Lists

To configure a SOCKS gateway, you can use the Management Console (easiest), the CLI, or you can create an installable list and load it on the ProxySG. To use the Management Console, see [Section A: "Configuring a SOCKS Gateway"](#) on page 840. For information on installing the file itself, see ["Creating a SOCKS Gateway Installable List"](#) on page 854.

The SOCKS gateways configuration includes SOCKS directives that:

- ❑ Names the SOCKS gateways, version, and port number
- ❑ Creates the SOCKS gateways groups
- ❑ Provide load balancing and host affinity
- ❑ Specifies the username
- ❑ Specifies the password

Available directives are described in the table below.

Table 41–1 SOCKS Directives

Directive	Meaning
gateway	Specifies the gateway alias and name, SOCKS port, version supported, usernames and password.
group	Creates a forwarding group directive and identifies member of the group.
host_affinity	Directs multiple connections by a single user to the same group member.
load_balance	Manages the load among SOCKS gateways in a group, or among multiple IP addresses of a gateway.
sequence alias_list	Adds a space-separated list of one or more SOCKS gateways and group aliases. (The default sequence is the default forwarding rule, used for all requests lacking policy instructions)
socks_fail	In case connections cannot be made, specifies whether to abort the connection attempt or to connect to the origin content server.

Syntax for the SOCKS directives are:

```
gateway gateway_alias gateway_name SOCKS_port [group=group_alias]
[version={4 | 5}] [user=username] [password=password] [encrypted-
password=encrypted_password]
group=group_alias [gateway_alias_list]
host_affinity http {none | client-ip-address | accelerator-cookie}
[gateway_or_group_alias]
host_affinity ssl {none | client-ip-address | accelerator-cookie |
ssl-session-id} [gateway_or_group_alias]
host_affinity other {none | client-ip-address}
[gateway_or_group_alias]
host_affinity timeout minutes
```

```

load_balance group {none | domain-hash | url-hash | round-robin |
least-connections} [group_alias]
load_balance gateway {none | round-robin | least-connections}
[gateway_alias]
sequence alias_list
socks_fail {open | closed}

```

For more information on SOCKS gateway directives, continue with the next section. For information on:

- ❑ `group` directives, continue with ["Creating SOCKS Gateways Groups Using Directives"](#) on page 851
- ❑ `load_balance` directives, continue with ["Configuring Load Balancing Directives"](#) on page 851
- ❑ `host_affinity` directives, continue with ["Configuring Host Affinity Directives"](#) on page 852
- ❑ `socks_fail` directives, continue with ["Setting Fail Open/Closed"](#) on page 851
- ❑ `sequence` directives, continue with ["Creating a Default Sequence"](#) on page 853

Configuring SOCKS Gateways Using Directives

SOCKS gateways can be configured using the gateways suboptions in the table below.

Table 41–2 SOCKS Gateways Syntax

Command	Suboptions	Description
gateway		Configures the SOCKS gateway.
	gateway_alias	A meaningful name that is used for policy rules.
	gateway_name	The IP address or name of the gateway where traffic is directed. The gateway name must DNS resolve.
	SOCKS_port	The port number of the SOCKS gateway.
	version={4 5}	The version that the SOCKS gateway can support.
	user=username	(Optional, if you use v5) The username of the user. It already must exist on the gateway.
	password=password	(Optional, if you use v5) The password of the user on the SOCKS gateway. It must match the gateway's information.
	encrypted-password=encrypted_password	(Optional, if you use v5) The encrypted password of the user on the SOCKS gateway. It must match the gateway's information.

Example

```
gateway Sec_Appl 10.25.36.47 1022 version=5 user=username
password=password
```

Creating SOCKS Gateways Groups Using Directives

The SOCKS gateway `groups` directive has the following syntax:

```
group group_name gateway_alias_1 gateway_alias_2...
```

where `group_name` is the name of the group, and `gateway_alias_1`, `gateway_alias_2`, and so forth are the gateways you are assigning to the SOCKS gateways group.

Setting Special Parameters

After you configure the SOCKS gateways and groups, you might need to set other special parameters to fine tune gateways. You can configure the following settings:

- ❑ "Setting Fail Open/Closed"
- ❑ "Configuring Load Balancing Directives" on page 851
- ❑ "Configuring Host Affinity Directives" on page 852

Setting Fail Open/Closed

Using directives, you can determine if the SOCKS gateways fails open or closed or if an operation does not succeed.

The syntax is:

```
socks_fail {open | closed}
```

where the value determines whether the SOCKS gateways should fail open or fail closed if an operation does not succeed. Fail open is a security risk, and fail closed is the default if no setting is specified. This setting can be overridden by policy, using the `SOCKS_gateway.fail_open(yes|no)` property.

Examples

```
socks_fail open
```

Configuring Load Balancing Directives

Load balancing shares the load among a set of IP addresses, whether a group or a gateway with multiple IP addresses.

The syntax is:

```
load_balance group {none | domain-hash | url-hash | round-robin |
least-connections} [group_alias]
load_balance gateway {none | round-robin | least-connections}
[gateway_alias]
```

Table 41–3 Load Balancing Directives

Command	Suboptions	Description
load_balance group	{none domain-hash url-hash round-robin least-connections} [group_alias]	If you use group for load balancing, you can set the suboption to none or choose another method. If you do not specify a group, the settings apply as the default for all groups.
load_balance gateway	{none round-robin least-connections} [gateway_alias]	If you use gateway for load balancing, you can set the suboption to none or choose another method. If you do not specify a gateway, the settings apply as the default for all gateways.

Example

```
load_balance gateway least_connections
```

Configuring Host Affinity Directives

Host affinity is the attempt to direct multiple connections by a single user to the same group member.

The syntax is:

```
host_affinity http {none | client-ip-address | accelerator-cookie}
[gateway_or_group_alias]
host_affinity ssl {none | client-ip-address | accelerator-cookie |
ssl-session-id} [gateway_or_group_alias]
host_affinity other {none | client-ip-address}
[gateway_or_group_alias]
host_affinity timeout minutes
```

Table 41–4 Commands to Configure Host Affinity Directives

Command	Suboption	Description
host_affinity http	{accelerator-cookie client-ip-address none} [gateway_or_group_alias]	Determines which HTTP host-affinity method to use (accelerator cookie or client-ip-address), or you can specify none. If you do not specify a gateway or group, the settings apply as the default for all gateways or groups.

Table 41–4 Commands to Configure Host Affinity Directives (Continued)

Command	Suboption	Description
host_affinity ssl	{accelerator-cookie client-ip-address none ssl-session-id} [gateway_or_group_alias]	Determines which SSL host-affinity method to use (accelerator-cookie, client-ip-address, or ssl-session-id), or you can specify none. If you do not specify a gateway or group, the settings apply as the default for all gateways or groups.
host_affinity other	other {none client-ip-address} [gateway_or_group_alias]	Determines whether TCP tunnel and Telnet is used. Determines whether to use the client-ip-address host-affinity method or specify none. If you do not specify a gateway or group, the settings apply as the default for all gateways or groups.
host_affinity timeout	minutes	Determines how long a user's IP address, SSL ID, or cookie remains valid when idle.

Example

```
host_affinity ssl accelerator-cookie 10.25.36.48
host_affinity timeout 5
```

Creating a Default Sequence

The default sequence is the default SOCKS gateways rule, used for all requests lacking policy instructions. Failover is supported if the sequence (only one is allowed) has more than one member.

Note: Set up sequences using policy. The default sequence (if present) is applied only if no applicable command is in policy.

For information on using VPM, refer to the *Visual Policy Manager Reference*; for information on using CPL, refer to the *Content Policy Language Guide*.

A default failover sequence works by allowing healthy SOCKS gateways to take over for an unhealthy gateway (one that is failing its DNS resolution or its health check). The sequence specifies the order of failover, with the second gateway taking over for the first gateway, the third taking over for the second, and so on). If all gateways are unhealthy, the operation fails either open or closed, depending upon your settings.

This configuration is generally created and managed through policy. If no forwarding policy applies, create a default sequence in the CPL or VPM.

Creating a SOCKS Gateway Installable List

You can create and install the SOCKS gateway installable list with the following methods:

- ❑ Use the Text Editor, which allows you to enter directives (or copy and paste the contents of an already-created file) directly onto the ProxySG.
- ❑ Create a local file on your local system; the ProxySG can browse to the file and install it.
- ❑ Use a remote URL, where you place an already-created file on an FTP or HTTP server to be downloaded to the ProxySG appliance.

When the SOCKS gateway installable list is created, it overwrites any previous SOCKS gateway configurations on the ProxySG. The installable list remains in effect until it is overwritten by another installable list; it can be modified or overwritten using Management Console or CLI commands.

Note: During the time that a SOCKS gateway installable list is being compiled and installed, SOCKS gateways might not be available. Any transactions that come into the appliance during this time might not be forwarded properly.

Installation of SOCKS gateway installable-list configuration should be done outside peak traffic times.

To create a SOCKS gateway installable list:

1. Select the **Configuration > Forwarding > SOCKS Gateways > Install SOCKS Gateway File** tab.
2. If you use a SOCKS gateway server for the primary or alternate forwarding gateway, you must specify the ID for the Identification (Ident) protocol used by the SOCKS gateway in SOCKS server handshakes. The default is `BLUECOAT SYSTEMS`.
3. From the drop-down list, select the method used to install the SOCKS gateway configuration; click **Install**.
 - **Remote URL:**

Enter the fully-qualified URL, including the filename, where the configuration is located. To view the file before installing it, click **View**. Click **Install**. Examine the installation status that displays; click **OK**.
 - **Local File:**

Click **Browse** to bring up the Local File Browse window. Browse for the file on the local system. Click **Install**. When the installation is complete, a results window opens. View the results, close the window, click **Close**.

- **Text Editor:**

The current configuration is displayed in installable list format. You can customize it or delete it and create your own. Click **Install**. When the installation is complete, a results window opens. View the results, close the window, click **Close**.

4. Click **Apply**.

Chapter 42: TCP Connection Forwarding

This section describes how to configure the ProxySG appliance to join peer clusters that process requests in asymmetrically routed networks.

Topics in this Section

The following topics are covered in this section:

- ["About Asymmetric Routing Environments"](#)
- ["The TCP Connection Forwarding Solution"](#) on page 858
- ["Configuring TCP Connection Forwarding"](#) on page 862

About Asymmetric Routing Environments

It is common in larger enterprises to have multiple ProxySG appliances residing on different network segments; for example, the enterprise receives Internet connectivity from more than one ISP. If IP spoofing is enabled, connection errors can occur because the ProxySG terminates client connections and makes a new outbound connection (with the source IP address of the client) to the server. The response might not return to the originating ProxySG, as illustrated in the following diagram.




Figure 42–1 Multiple ProxySG appliances in an asymmetric routing environment

After a connection occurs (either intercepted or bypassed) through any ProxySG in the connection forwarding cluster, future packets of any such recorded flow that is subject to asymmetric routing are properly handled. The ProxySG also recognizes self-originated traffic (from any of the peers of the connection forwarding cluster), so any abnormal internal routing loops are also appropriately processed.

The TCP Connection Forwarding Solution

Enabling TCP Connection Forwarding is a critical component of the following solutions:

- ❑ "About Bidirectional Asymmetric Routing" on page 859.
- ❑ "About Dynamic Load Balancing" on page 860.
- ❑ "About ADN Transparent Tunnel Load Balancing" on page 861.

About Bidirectional Asymmetric Routing

To solve the asymmetric routing problem, at least one ProxySG on each network segment must be configured to perform the functionality of an L4 switch. These selected appliances form a cluster. With this peering relationship, the connection responses are able to be routed to the network segment where the originating client resides.

Starting in the SGOS 5.1.4.x release, cluster membership is manual; that is, ProxySG appliances must be added to a cluster by enabling connection forwarding and adding a list of other peers in the cluster. After a peer joins a cluster, it begins sending and receiving TCP connections, and notifies the other peers about its connection requests.




Figure 42–2 ProxySG appliances share TCP connection information

About Dynamic Load Balancing

In a deployment where one ProxySG receives all of the traffic originating from clients and servers from an external routing device and distributes connections to other ProxySG appliances, TCP connection forwarding enables all of the appliances to share connection information (for each new connection) and the in-line ProxySG routes the request back to the originating appliance, thus lightening the load on the in-path appliance.




Figure 42-3 A ProxySG appliance serving in-path as a load balancer

In the above network topography, ProxySG **SG 1** is deployed in-path to receive all traffic (by way of a switch) originating from the clients to the servers and servers to the clients and serves as a load balancer to the other four ProxySG appliances. Appliances **2** through **5** also have independent connectivity to the clients and the servers. When all appliances belong to the same peering cluster and have connection forwarding enabled, appliance **SG 1** knows which of the other appliances made a specific connection and routes the response to that appliance.

In this deployment, a TCP acknowledgement is sent and retransmitted, if required, to ensure the information gets there, but each new connection message is not explicitly acknowledged. However, if the ProxySG receives packets for a connection that is unrecognized, the appliance retains those packets for a short time before deciding whether to forward or drop them, which allows time for a new connection message from a peer to arrive.

While adding more peers to a cluster increases the connection synchronization traffic, the added processing power all but negates that increase. You can have multiple peer clusters, and if you are cognoscente of traffic patterns to and from each cluster, you can create an effective cluster strategy. The only limitation is that a ProxySG can only be a peer in one cluster.

The Blue Coat load balancing solution is discussed in greater detail in earlier sections.

About ADN Transparent Tunnel Load Balancing

TCP connection forwarding is a critical component of the Blue Coat ADN transparent tunnel load balancing deployment. Achieving efficient load balancing is difficult when ADN transparent tunneling is employed and an external load balancer is distributing requests to multiple ProxySG appliances.

A user-noticeable performance degradation occurs if the router, switch, or load balancer sends traffic to a ProxySG that has not been servicing a particular client long enough to build up substantial byte caching dictionary, thus the compression ratio is low. When the ProxySG appliances connected to the routing device belong to the same peer cluster and connection forwarding is enabled, the ADN managers on each appliance know which of their peers has the best byte caching dictionary with the client and forwards the request. This is illustrated in the following diagram.




Figure 42–4 ADN Transparent Tunnel load balancing with Connection Forwarding enabled

Load balancing is based on the IP address of the remote ADN peer. This assures that all the traffic from a particular ADN peer to the local ADN cluster always goes to a specific local ProxySG, thus eliminating the inefficiency of keeping dictionaries for that remote peer on more than one local ProxySG.

The Blue Coat ADN solution is discussed in greater detail in ["Configuring an Application Delivery Network"](#) on page 711.

TCP Configuration Forwarding Deployment Notes

When configuring your network for TCP connection forwarding, consider the following:


- ❑ Peers can be added to clusters at any time without affecting the performance of the other peers. A ProxySG that joins a peer cluster immediately contacts every other peer in the cluster. Likewise, a peer can leave a cluster at anytime. This might be a manual drop or a forced drop because of a hardware or software failure. If this happens, the other peers in the cluster continue to process connection forwarding requests.
- ❑ Connections between peers are not encrypted and not authenticated. If you do not assign the correct local IP address on a ProxySG with multiple IP addresses, traffic sent peer to peer might be routed through the Internet, not the intranet, exposing your company-sensitive data.
- ❑ The peering port—the connection between ProxySG connection forwarding peers—cannot be configured with bypass services. This means a ProxySG cannot be deployed in transparent mode between two ProxySG appliances that are peers.
- ❑ The ProxySG does not enforce a maximum number of appliances a peer cluster supports, but currently the deployment is designed to function with up to 20 ProxySG appliances.
- ❑ Because TCP connection forwarding must function across different network segments, employing multicasting, even among ProxySG peers on the same network, is not supported.
- ❑ There might be a slight overall performance impact from enabling TCP connection forwarding, especially in deployments where traffic is largely already being routed to the correct ProxySG. If a substantial amount of traffic requires forwarding, the performance hit is equitable to processing the same amount of bridging traffic.

Configuring TCP Connection Forwarding

As described in the previous concept sections, enabling TCP connection forwarding provides one component to a larger deployment solution. After you have deployed Blue Coat appliances into the network topography that best fits your enterprise requirements, enable TCP connection forwarding on each Blue Coat appliance that is to belong to the peering cluster, and add the IP address of the other peers. The peer lists on *all* of the cluster members must be the same, and a ProxySG cannot have a different local peer IP address than what is listed in another peers list. A peer list can contain only one local IP address.

To enable TCP Connection Forwarding:

1. Select the **Configuration > Network > Advanced > Connection Forwarding** tab.



2. From the **Local IP** drop-down list, select the IP address that is routing traffic to this ProxySG.

Specify the port number (the default is **3030**) that the ProxySG uses to communicate with all peers, which includes listening and sending out connection forwarding cluster control messages to all peers in the group. *All* peers in the group must use the same port number (when connection forwarding is enabled, you cannot change the port number).

3. Add the cluster peers:
 - a. Click **Add**.
 - b. In the **Peer IPs** field, enter the IP addresses of the other peers in the cluster that this ProxySG is to communicate connection requests with; click **OK**.
4. Select **Enable Connection Forwarding**.
5. Click **Apply**.

This ProxySG joins the peer cluster and immediately begins communicating with its peers.

Copying Peers to Another ProxySG in the Cluster

If you have a larger cluster that contains several peer IP addresses, select all of the IP addresses in the **Connection Forwarding Peer IPs** list and click **Copy To Clipboard**; this action includes the local IP address of the peer you are copying from, and it will be correctly added as a remote peer IP address on the next appliance. When you configure connection forwarding on the next appliance, click **Paste From**

Clipboard to paste the list of peers, and click **Apply**. Whichever peer IP address is the new appliance's local IP address is pulled out of the list and used as the local IP address on the new appliance. If a local IP address is not found or if more than one local IP address is found, the paste fails with an error.

Removing a Peer

A network change or other event might require you to remove a peer from the cluster. Highlight a peer IP address and click **Remove**. The peer connection is terminated and all connections associated with the peer are removed from the local system.

You can also remove all peers from the cluster by clicking the **Remove...** button. A dialog appears, asking you to confirm your choice to remove all peers.

Chapter 43: Configuring the Upstream Network Environment

The following topics in this chapter discuss how to configure the ProxySG to interact with both the local network and with the upstream network environment:

- ❑ [Section A: "Overview" on page 866](#)
- ❑ [Section B: "About Forwarding" on page 867](#)
- ❑ [Section C: "Configuring Forwarding" on page 875](#)
- ❑ [Section D: "Using Forwarding Directives to Create an Installable List" on page 883](#)

Section A: Overview

To control upstream interaction, the ProxySG supports the following:

- ❑ The ProxySG forwarding system—Allows you to define the hosts and groups of hosts to which client requests can be redirected. Those hosts can be servers or proxies. Rules to redirect requests are set up in policy.
- ❑ SOCKS gateways—SOCKS servers provide application-level firewall protection for an enterprise. The SOCKS protocol provides a generic way to proxy HTTP and other protocols. For information on configuring SOCKS gateways, see [Chapter 41: "SOCKS Gateway Configuration"](#) on page 839.

Section B: About Forwarding

Forwarding creates a *proxy hierarchy*, which consists of a set of proxies (including ProxySG appliances that are configured as proxies (**Configuration > Proxy Services**)). Appliances close to the origin server perform object caching for server content and distribute the content to the object caches of other proxies that are farther away from the origin server. If forwarding is set up in an organized manner, the load involved with object caching is distributed throughout the proxy hierarchy, which avoids sending any piece of content across any given WAN link more than once.

For more information, see one of the following topics:

- ❑ ["About the Forwarding System"](#)
- ❑ ["Example of Using Forwarding"](#) on page 867
- ❑ ["About Load Balancing and Health Checks"](#) on page 872
- ❑ ["About Host Affinity"](#) on page 873
- ❑ ["Using Load Balancing with Host Affinity"](#) on page 874

To get started configuring forwarding, see [Section C: "Configuring Forwarding"](#) on page 875.

About the Forwarding System

Forwarding redirects content requests to IP addresses other than those specified in the requesting URL. Forwarding affects only the IP address of the upstream device to which a request is sent; forwarding does *not* affect the URL in the request.

The ProxySG forwarding system consists of forwarding, upstream SOCKS gateways, load balancing, host affinity, and health checks. The forwarding system determines the upstream address where a request is sent, and is tied in with all the protocol proxies.

Note: The ProxySG forwarding system directly supports the forwarding of HTTP, HTTPS, FTP, MMS, RTSP, Telnet, and TCP tunnels.

For more information, see one of the following topics:

- ❑ ["Example of Using Forwarding"](#)
- ❑ ["About Load Balancing and Health Checks"](#) on page 872
- ❑ ["About Host Affinity"](#) on page 873
- ❑ ["Using Load Balancing with Host Affinity"](#) on page 874
- ❑ [Section C: "Configuring Forwarding"](#) on page 875

Example of Using Forwarding

This section discusses an example of using forwarding to minimize traffic over WAN links and to the Internet by leveraging object caching on proxies in the forwarding system.

For more information, see the following topics:

- ❑ ["High-Level View of the Example System"](#)
- ❑ ["Example Network" on page 869](#)
- ❑ ["How the Example Uses Object Caching" on page 870](#)

See Also

["About the Forwarding System" on page 867](#)

High-Level View of the Example System

The example discussed in this section uses the following logical proxy hierarchy.




Figure 43-1 Logical proxy hierarchy used in the forwarding example

In [Figure 43-1](#), there are five ProxySG appliances configured as proxies: one in the central data center and one apiece in four branch offices or sites. ProxySG 1, located in the central data center, provides Internet access for the entire system.

ProxySG 4 uses ProxySG 2 as its forwarding host, and ProxySG 2 uses ProxySG 1 as its forwarding host. Similarly, ProxySG 5 uses ProxySG 3 as its forwarding host and ProxySG 3 uses ProxySG 1 as its forwarding host.

This means that, for example, any piece of content in ProxySG 1's object cache can be distributed to ProxySG 2 or ProxySG 3's object cache without having to send the content over the Internet.

Continue with "Example Network" .

Example Network

The following figure shows a more detailed view of the example network.




Figure 43-2 Example ADN network that uses forwarding

In Figure 43-2, ProxySG 1 (located in the central data center) acts as a gateway to the Internet; in other words, all Internet access goes through ProxySG 1. Two regional data centers accept requests from four branch offices or sites, each with ProxySG appliances configured as proxies.

Because ProxySG 1 is the gateway to the Internet, it is *upstream* of all other ProxySGs shown in [Figure 43–2](#). Because load-balanced ProxySG 2 and ProxySG 3 are configured to use ProxySG 1 as their forwarding host, they are *downstream* of ProxySG 1. And because ProxySG 4 is configured to use the load-balanced group of ProxySGs labeled 2 as its forwarding host, ProxySG 4 is downstream of both ProxySG 2 and ProxySG 1. Finally, ProxySG 5 is downstream of ProxySG 3 and ProxySG 1.

Another way of stating this, using ProxySG 4 as an example, is that any request to the Internet goes through ProxySG 2 and then ProxySG 1 instead of going directly to the host specified in the URL of the request.

Continue with "[How the Example Uses Object Caching](#)".

How the Example Uses Object Caching

[Figure 43–3](#) shows how forwarding and object caching work together to minimize traffic over the example network's WAN links and to the Internet.




Figure 43–3 How forwarding can leverage object caching to prevent multiple requests to the Internet and over WAN links

In [Figure 43–3](#), a user connected to ProxySG 4 requests content located on a Web server in the Internet. The content—which might be a spreadsheet or multimedia—is in the object cache of load-balanced ProxySG 2, and therefore is retrieved from the object cache. Neither the WAN links nor the origin server are used to retrieve the content. The content is then cached on ProxySG 4’s object cache so the next time a user requests the same content, it is retrieved from ProxySG 4’s object cache.

If a user connected to ProxySG 5 requests the same content—and the content is in neither ProxySG 5’s nor ProxySG 3’s object cache—load-balanced ProxySG 3 gets the content from ProxySG 1 and object caches it. Subsequently, ProxySG 5 gets the content from ProxySG 3 and object caches it.

Because the content is in ProxySG 1's object cache, the content is not retrieved from the origin server. In this scenario, only the WAN links are used; the Internet link is *not* used to retrieve the content.

Note: In [Figure 43–2](#) and [Figure 43–3](#), each ProxySG is assumed to use one IP address for forwarding. You could achieve similar results using load balancing if you configure a DNS host name as a forwarding host and used DNS load balancing to forward requests to more than one ProxySG. For more information, see ["About Load Balancing and Health Checks"](#) .

See Also

["About Load Balancing and Health Checks"](#)

["About Host Affinity"](#) on page 873

["About the Forwarding System"](#) on page 867

About Load Balancing and Health Checks

Load balancing distributes forwarding traffic among multiple IP addresses to achieve optimal resource utilization, to maximize throughput, and to minimize response time. Typically, you use load balancing to distribute requests to more than one ProxySG appliance, although you can also distribute requests to multiple IP addresses on a single appliance—or a combination of the two.

This section discusses the following topics:

- ❑ ["Load Balancing Methods"](#)
- ❑ ["Health Checks"](#) on page 873

Load Balancing Methods

ProxySG load balancing methods include *round robin*—which selects the next system in the list—or *least connections*—which selects the system with the fewest number of connections.

You can configure load balancing in any of the following ways:

- ❑ For individual hosts: If a host is DNS-resolved to multiple IP addresses, then that host's load-balancing method (round robin, least connections, or none) is applied to those IP addresses. The method is either explicitly set for that host or taken from the configurable global default settings.
- ❑ For groups: Load balancing for groups works exactly the same as load balancing for hosts with multiple IP addresses except there are two additional load balancing methods for groups:
 - URL hash—Requests are hashed based on the request URL.
 - Domain hash—Requests are hashed based on the domain name in the request.

Continue with ["Health Checks"](#) .

Health Checks

The availability of a proxy to participate in load balancing depends on the status of the proxy's health check (**Statistics > Health Checks**). The name of a forwarding hosts or group starts with `fwd.`; any host or group whose health status is Unhealthy is excluded from forwarding.

If a proxy has a health check of Unhealthy, the proxy is assumed to be down and cannot participate in load balancing. If this happens, verify the following:

- ❑ The proxy or proxies are all intercepting traffic on the same ports you configured in your forwarding host or group.

If the health check for a downstream proxy is shown as unhealthy on the upstream proxy, verify that the downstream proxy intercepts traffic on the specified port in the forwarding host on the upstream proxy.

For example, if you set up forwarding for HTTP traffic on port 80, make sure the forwarding proxy or proxies are set to intercept HTTP traffic on port 80 (**Services > Proxy Services**).

- ❑ The proxy or proxies are available. Use the `ping` command from a downstream proxy to verify upstream proxies are available.
- ❑ Verify the proxies' health status and take corrective action if necessary.

For more information, see [Chapter 73: "Verifying the Health of Services Configured on the ProxySG"](#) on page 1353.

In the event no load balancing host is available, *global defaults* determine whether the connection fails open (that is, goes directly to its destination) or fails closed (that is, the connection fails). For more information, see ["Configuring Global Forwarding Defaults"](#) on page 879.

About Host Affinity

Host affinity is the attempt to direct multiple connections by a single user to the same group member. Host affinity causes the user's connections to return to the same server until the configurable host affinity timeout period is exceeded.

For example, suppose a Web site with a shopping cart has several load-balanced Web servers, but only one Web server has the session data for a given user's shopping cart transaction. If a connection is sent to a different Web server that has no data about the user's session, the user has to start over. ProxySG host affinity helps make sure each request goes to its proper destination; however, the proxy does *not* interact with the session or with session data.

Host affinity allows you to use the following options:

- ❑ Use the client IP address to determine which group member was last used. When the same client IP sends another request, the host makes the connection to that group member.
- ❑ Place a cookie in the response to the client. When the client makes future requests, the cookie data is used to determine which group member the client last used. The host makes the connection to that group member.

- ❑ For HTTPS, extract the SSL session ID name from the connection information. The host uses the session ID in place of a cookie or client IP address to determine which group member was last used. The host makes the connection to that group member.

Using Load Balancing with Host Affinity

Symantec highly recommends that if you enable load balancing, you also enable host affinity.

By default, if you use load balancing, each connection is treated independently. The connection is made to whichever member of the load-balancing group the load-balancing algorithm selects.

If host affinity is configured, the system checks host affinity first to see if the request comes from a known client. If this is a first connection, the load-balancing algorithm selects the group member to make the connection. Host affinity records the result of the load balancing and uses it if that client connects again.

Host affinity does not make a connection to a host that health checks report is down; instead, if host affinity breaks, the load-balancing algorithm selects a group member that is healthy and re-establishes affinity on that working group member.

Host affinity methods are discussed in the following table.

Table 43–1 Host Affinity Methods

Setting	Description	HTTP	SSL	Other (TCP Tunnel or Telnet)
Global Default	Use the default setting for all forwarding hosts on the system.	x	x	x
None	Disables host affinity.	x	x	x
Client IP Address	Uses the client IP address to determine which forwarding group member was last used.	x	x	x
Accelerator Cookie	Inserts a cookie into the response to the client.	x	x	
SSL Session ID	Used in place of a cookie or client IP address. Extracts the SSL session ID name from the connection information.		x	

Section C: Configuring Forwarding

High-level steps to configure forwarding are:

- ❑ Create the forwarding hosts and groups, including parameters such as protocol agent and port.
- ❑ Set load balancing and host affinity values.

See Also

["Creating Forwarding Hosts and Groups"](#) on page 875

["About the Forwarding System"](#) on page 867

["Example of Using Forwarding"](#) on page 867

Creating Forwarding Hosts and Groups

Before you can create forwarding groups, you must create forwarding hosts as discussed in this section. A forwarding host is a ProxySG configured as a proxy to which certain traffic is redirected for the purpose of leveraging object caching to minimize trips to the Internet and over WAN links.

For more information about forwarding hosts, see [Section B: "About Forwarding"](#) on page 867.

This section discusses the following topics:

- ❑ ["Creating Forwarding Hosts"](#)
- ❑ ["Creating Forwarding Groups"](#) on page 877

You can create as many hosts or groups as you need.

Creating Forwarding Hosts

This section discusses how to create a forwarding host. To create a forwarding group, see ["Creating Forwarding Groups"](#) on page 877.

To create forwarding hosts:

1. Select the **Configuration > Forwarding > Forwarding Hosts** tab.
2. Click **New**. The Add Forwarding Host dialog displays.

3. Configure the host options:

- a. In the **Alias** field, enter a unique name to identify the forwarding host.

Note: Because the forwarding host alias is used in policy, the alias cannot be a CPL keyword, such as `no`, `default`, or `forward`.

- b. In the **Host** field, enter the forwarding host's fully qualified domain name or IPv4/IPv6 address.

- c. For **Type**, click one of the following:

- **Server** should be used for reverse proxy deployments. Choosing **Server** means you will use the relative path for URLs in the HTTP header because the next hop is a Web server, not a proxy server. HTTPS, TCP tunnels, and Telnet can be forwarded to a server only; they cannot be forwarded to a proxy.
- **Proxy** should be used in forward proxy deployments.

- d. Select the option next to each protocol to forward.

In the adjacent **Port** field, enter the port you want to use for forwarding. Port 80 is the default for HTTP. The rest of the host types default to their appropriate Internet default port, except TCP tunnels, which have no default and for which a port must be specified.

- e. In the **Load Balancing and Host Affinity** section, make the following selections:
 - From the **Load balancing method** list, click one of the following:
 - **Global default** (configured on the **Configuration > Forwarding > Global Defaults** tab), which sets the default for all forwarding hosts on the system.
 - **Round Robin**, which causes the request to be forwarded to the next forwarding host or group in the sequence.
 - **Least Connections**, which causes requests to be sent to the forwarding host or group that currently has the least number of connections.
 - **None**, which means load balancing will not be used.
 - From the **Host affinity methods** list (see [Table 43–1, "Host Affinity Methods"](#)), click the method you want to use.
4. Click **OK**.
5. Click **Apply**.

See Also

["Creating Forwarding Groups"](#)

Section D: ["Using Forwarding Directives to Create an Installable List"](#) on page 883

["About the Forwarding System"](#) on page 867

["Example of Using Forwarding"](#) on page 867

Creating Forwarding Groups

This section discusses how to create a forwarding group. To create a forwarding host, see ["Creating Forwarding Hosts"](#) on page 875.

To create forwarding groups:

An existing host can belong to one or more groups as needed. It can belong only once to a single group.

1. Select the **Configuration > Forwarding > Forwarding Groups** tab.
2. Click **New**. The Add Forwarding Group dialog displays, showing the available aliases.

3. In the **Alias** field, enter a unique name to identify the forwarding group.

Note: Because the forwarding group alias is used in policy, the alias cannot be a CPL keyword, such as `no`, `default`, or `forward`.

4. To add members to a group, click the name of the hosts you want grouped and click **Add**.
5. Choose load balancing and host affinity methods:
 - From the **Load balancing method** list, click one of the following:
 - **Global default** (configured on the **Configuration > Forwarding > Global Defaults** tab), which sets the default for all forwarding hosts on the system.
 - **Round Robin**, which causes the request to be forwarded to the next forwarding host or group in the sequence.
 - **Least Connections**, which causes requests to be sent to the forwarding host or group that currently has the least number of connections.
 - **Url Hash**, which hashes requests based on the request URL.

- **Domain Hash**, which hashes requests based on the domain name in the request.
 - **None**, which means load balancing will not be used.
 - From the **Host affinity methods** list (see [Table 43–1, "Host Affinity Methods"](#)), click the method you want to use.
6. Click **OK**.
 7. Click **Apply**.

See Also

["Creating Forwarding Hosts"](#)

[Section D: "Using Forwarding Directives to Create an Installable List" on page 883](#)

["About the Forwarding System" on page 867](#)

["Example of Using Forwarding" on page 867](#)

Configuring Global Forwarding Defaults

The global defaults apply to all forwarding hosts and groups that are configured for **Use Global Default**. For example, if you choose **Use Global Default** for **Load Balancing Method** in the definition of a forwarding host or group, this section discusses how to configure those default settings.

To configure global defaults:

1. Select the **Configuration > Forwarding > Global Defaults** tab.

The screenshot shows the 'Global Defaults' configuration page. At the top, there are three tabs: 'Forwarding Hosts', 'Forwarding Groups', and 'Global Defaults', with 'Global Defaults' being the active tab. The page is divided into two main sections: 'General Settings' and 'Load Balancing and Host Affinity'.

General Settings:

- If no forwarding host is available:**
 - ☐ Connect directly (fail open)
 - ☒ Deny the request (fail closed)
- ☒ Use forwarding for administrative downloads
- Timeout for integrated hosts: minutes

Load Balancing and Host Affinity:

- Load balancing methods:**
 - Forwarding hosts:
 - Forwarding groups:
- Host affinity methods:**
 - HTTP:
 - SSL:
 - Other:
- Host affinity timeout: minutes

2. Configure the **General Settings**:

- a. Determine how connections behave if no forwarding is available. Failing open is an insecure option. The default is to fail closed. This setting can be overridden by policy, if it exists.
- b. Decide if you want to **Use forwarding for administrative downloads**. The default is to use forwarding in this case.

This option determines whether forwarding is applied to requests generated for administrative reasons on the system, such as downloading policy files or new system images.

If the option is on, meaning that forwarding is applied, you can control the forwarding in policy as needed.

This option also affects the use of SOCKS gateways.

- c. Enter the **Timeout for integrated hosts** interval: An integrated host is an Origin Content Server (OCS) that has been added to the health check list. The host, added through the `integrate_new_hosts` policy property, ages out after being idle for the specified time. The default is 60 minutes.

3. Configure **Global Load Balancing** and **Host Affinity Settings**.

- a. Load-balancing methods:
 - Forwarding hosts: Specify the load-balancing method for all forwarding hosts unless their configuration specifically overwrites the global settings. You can choose **Least Connections** or **Round Robin**, or you can disable load balancing by selecting **None**. **Round Robin** is specified by default.
 - Forwarding groups: Specify the load-balancing method for all forwarding groups unless their configuration specifically overwrites the global settings. You can choose to do a **domain hash** or a **URL hash**. You can also select **Least Connections** or **Round Robin**, or disable load balancing by selecting **None**. **Round Robin** is specified by default.
- b. In the **Global Host Affinity** methods area (see [Table 43–1, "Host Affinity Methods"](#)), select the method you want to use.
- c. Enter the **Host Affinity Timeout** interval, the amount of time a user's IP address, SSL ID, or cookie remains valid after its most recent use. The default is 30 minutes, meaning that the IP address, SSL ID or cookie must be used once every 30 minutes to restart the timeout period.

4. Click **Apply**.

Configuring the Forwarding Default Sequence

The default sequence is the forwarding sequence used when there is no matching forwarding rule in policy.

Following is an example of forwarding policy:

```
<Forward>
url.domain=bluecoat.com forward(FWGrp2, FWGrp1)
```

In the example, requests that match the URL domain `bluecoat.com` are sent to a forwarding group named `FWGrp2` unless all of the members in `FWGrp2` are down, in which case requests are sent to `FWGrp1`. Health checks are performed continually to minimize the possibility that requests are sent to a forwarding host or group that is known to be down.

The default sequence (and any sequence specified in policy) works by allowing healthy hosts to take over for an unhealthy host or group (one that is failing its DNS resolution or its health check). If more than one member is in the sequence, the sequence specifies the order of failover, with the second host or group taking over for the first one, the third taking over for the second, and so on.

If all of the hosts in the sequence are down, the request either fails open or fails closed (that is, the connection is denied). Symantec recommends you set this behavior in policy as follows:

```
forward.fail_open(yes|no)
```


However, you can also configure it using global defaults as discussed in ["Configuring Global Forwarding Defaults"](#) on page 879.

Note: The CLI command `#(config forwarding)sequence {add alias-name | clear | demote alias-name | promote alias-name | remove alias-name}` is intended for backward compatibility with previous SGOS versions for which there is no equivalent CPL. Symantec recommends that you create forwarding policy (including sequences) using CPL or VPM.

For information on using VPM, refer to the *Visual Policy Manager Reference*; for information on using CPL, refer to the *Content Policy Language Guide*. For information on using forwarding with policy, see [Chapter 44: "Using Policy to Manage Forwarding"](#) on page 891.

To create the default sequence:

1. Select the **Configuration > Forwarding > Default Sequence** tab. The available aliases display.



2. To select an alias, click its name in the **Available Aliases** area and click **Add**.

Note: Any host or group in the default sequence is considered in use by policy. As a result, if you try to delete a host or group while it is in the default sequence, you receive an error message. You must remove the host/group from the sequence first, then delete the host or group.

3. Click **Promote** or **Demote** to change the order of the hosts in the default sequence.
4. Click **Apply**.

Statistics

To view forwarding statistics, select the **Statistics > Advanced > Forwarding** tab.

Section D: Using Forwarding Directives to Create an Installable List

The information in this section is provided for backward compatibility only.

You can use directives instead of using the Management Console or CLI to configure forwarding. Using directives, you can:

- ❑ Create the forwarding hosts and groups
- ❑ Provide load balancing and host affinity

This section discusses the following topics:

- ❑ ["Creating Forwarding Host and Group Directives"](#)
- ❑ ["Setting Special Parameters"](#) on page 886
- ❑ ["Creating a Forwarding Default Sequence"](#) on page 888
- ❑ ["Creating a Forwarding Installable List"](#) on page 889

Table 43–2 Forwarding Directives

Directive	Meaning	See
fwd_fail	Determines whether the forwarding host should fail open or fail closed if an operation does not succeed.	"Setting Fail Open/Closed and Host Timeout Values" on page 886.
fwd_host	Creates a forwarding host and sets configuration parameters for it, including protocols and ports.	"Creating Forwarding Hosts Using Directives" on page 884.
group	Creates a forwarding group and identifies members of the group.	"Creating Forwarding Groups Using Directives" on page 885.
host_affinity	Directs multiple connections by a single user to the same group member.	"Configuring Host Affinity Directives" on page 887.
integrated_host_timeout	Manages an origin content server that has been added to the health check list. The host ages out after being idle for the specified time.	"Setting Fail Open/Closed and Host Timeout Values" on page 886.
load_balance	Manages the load among forwarding hosts in a group, or among multiple IP addresses of a host.	"Configuring Load-Balancing Directives" on page 887.

Table 43–2 Forwarding Directives (Continued)

Directive	Meaning	See
sequence	Sets the default sequence to the space separated list of one or more forwarding host and group aliases. (The default sequence is the default forwarding rule, used for all requests lacking policy instructions.)	"Creating a Forwarding Default Sequence" on page 888.

Creating Forwarding Host and Group Directives

A forwarding host directive creates a host along with all its parameters. You can include a group that the forwarding host belongs to.

A group directive creates a group and identifies group members.

This section discusses the following topics:

- "Creating Forwarding Hosts Using Directives"
- "Creating Forwarding Groups Using Directives" on page 885

Creating Forwarding Hosts Using Directives

To create a forwarding host, choose the protocols you want to use and add the forwarding host to a group, enter the following into your installable list. Create a `fwd_host` directive for each forwarding host you want to create.

```
fwd_host host_alias hostname [http[=port]] [https[=port]] [ftp[=port]]
[mms[=port]] [rtsp[=port]] [tcp=port] [telnet[=port]] [ssl-verify-
server[=yes | =no]] [group=group_name [server | proxy]]
```

Table 43–3 Commands to Create Forwarding Host and Group Directives

host_alias		This is the alias for use in policy. Define a meaningful name.
hostname		The name of the host domain, such <code>www.bluecoat.com</code> , or its IP address.
http https ftp mms rtsp telnet	=port	At least one protocol must be selected. HTTPS and Telnet cannot be used with a proxy. Note that HTTPS refers to terminated HTTPS, so it is used only for a server.
tcp	=port	If you choose to add a TCP protocol, a TCP port must be specified. TCP protocols are not allowed if the host is a proxy.

Table 43–3 Commands to Create Forwarding Host and Group Directives (Continued)

<code>ssl-verify-server</code>	<code>=yes =no</code>	Sets SSL to specify that the ProxySG checks the CA certificate of the upstream server. The default for <code>ssl-verify-server</code> is <code>yes</code>. This can be overridden in the SSL layer in policy. To disable this feature, you must specify <code>ssl-verify-server=no</code> in the installable list or CLI. In other words, you can configure <code>ssl-verify-server=yes</code> in three ways: do nothing (<code>yes</code> is the default), specify <code>ssl-verify-server=no</code>, or specify <code>ssl-verify-server=yes</code>.
<code>group</code>	<code>=group_name</code>	Specifies the group (or server farm or group of proxies) to which this host belongs. If this is the first mention of the group <code>group_name</code> then that group is automatically created with this host as its first member. The ProxySG uses load balancing to evenly distribute forwarding requests to the origin servers or group of proxies.
<code>server proxy</code>		<code>server</code> specifies to use the relative path for URLs in the HTTP header because the next hop is a Web server, not a proxy server. The default is <code>proxy</code> .

Example

```
fwd_host www.bluecoat1.com 10.25.36.48 ssl-verify-server=no
group=bluecoat
```

See Also

["Creating Forwarding Groups Using Directives"](#)

Creating Forwarding Groups Using Directives

The forwarding groups directive has the following syntax:

```
group group_name host_alias_1 host_alias_2...
```

where `group_name` is the name of the group, and `host_alias_1`, `host_alias_2`, and so forth are the forwarding hosts you are assigning to the forwarding group.

Forwarding host parameters are configured through the forwarding host directives.

See Also

["Creating Forwarding Hosts Using Directives"](#) on page 884

Setting Special Parameters

After you configure the forwarding hosts and groups, you might need to set other special parameters to fine tune the hosts. You can configure the following settings:

- ❑ ["Setting Fail Open/Closed and Host Timeout Values"](#)
- ❑ ["Configuring Load-Balancing Directives"](#) on page 887
- ❑ ["Configuring Host Affinity Directives"](#) on page 887

Setting Fail Open/Closed and Host Timeout Values

Using directives, you can determine if the forwarding host fails open or closed, if an operation does not succeed, and the interval it takes for integrated hosts to be aged out.

An integrated host is an Origin Content Server (OCS) that has been added to the health check list. If the policy property `integrate_new_hosts` applies to a forwarding request as a result of matching the `integrate_new_hosts` property, the ProxySG makes a note of each OCS and starts health checking to help future accesses to those systems. If the host is idle for the interval you specify, it is aged out. Sixty minutes is the default interval.

The syntax is:

```
fwd_fail {open | closed}
integrated_host_timeout minutes
```

Table 43–4 Commands to Set Fail Open/Closed and Host Timeout Values

<code>fwd_fail</code>	<code>{open closed}</code>	Determines whether the forwarding host should fail open or fail closed if an operation does not succeed. Fail open is a security risk, and fail closed is the default if no setting is specified. This setting can be overridden by policy, (using the <code>forward.fail_open(yes no)</code> property).
<code>integrated_host_timeout</code>	<code>minutes</code>	An OCS that has been added to the health check list is called an integrated host. The host ages out after being idle for the specified time.

Examples

```
fwd_fail open
integrated_host_timeout 90
```

See Also

["Configuring Load-Balancing Directives"](#)

["Configuring Host Affinity Directives"](#) on page 887

Configuring Load-Balancing Directives

Load balancing shares the load among a set of IP addresses, whether a group or a host with multiple IP addresses.

The syntax is:

```
load_balance group {none | domain-hash | url-hash | round-robin |
least-connections} [group_alias]
load_balance host {none | round-robin | least-connections}
[host_alias]
```

Table 43–5 Load Balancing Directives

Command	Suboptions	Description
load_balance group	{none domain-hash url-hash round-robin least-connections} [group_alias]	If you use group for load balancing, you can set the suboption to none or choose another method. If you do not specify a group, the settings apply as the default for all groups.
load_balance host	{none round-robin least-connections} [host_alias]	If you use host for load balancing, you can set the suboption to none or choose another method. If you do not specify a host, the settings apply as the default for all hosts.

Example

```
load_balance host least_connections
```

See Also

["Configuring Host Affinity Directives"](#)

["Creating a Forwarding Default Sequence" on page 888](#)

["Creating a Forwarding Installable List" on page 889](#)

Configuring Host Affinity Directives

Host affinity is the attempt to direct multiple connections by a single user to the same group member.

The syntax is:

```
host_affinity http {none | client-ip-address | accelerator-cookie}
[host_or_group_alias]
host_affinity ssl {none | client-ip-address | accelerator-cookie |
ssl-session-id} [host_or_group_alias]
host_affinity other {none | client-ip-address} [host_or_group_alias]
host_affinity timeout minutes
```

Table 43–6 Commands to Configure Host Affinity Directives

Command	Suboption	Description
host_affinity http	{accelerator-cookie client-ip-address none} [host_or_group_alias]	Determines which HTTP host-affinity method to use (accelerator cookie or client-ip-address), or you can specify none. If you do not specify a host or group, the settings apply as the default for all hosts or groups.
host_affinity ssl	{accelerator-cookie client-ip-address none ssl-session-id} [host_or_group_alias]	Determines which SSL host-affinity method to use (accelerator cookie, client-ip-address, or ssl-session-id), or you can specify none. If you do not specify a host or group, the settings apply as the default for all hosts or groups.
host_affinity other	{none client-ip- address} [host_or_group_alias]	Determines whether client-ip-address mode is used with TCP tunnels or Telnet.
host_affinity timeout	minutes	Determines how long a user's IP address, SSL ID, or cookie remains valid when idle

Example

```
host_affinity ssl_method 10.25.36.48
host_affinity timeout 5
```

See Also

"Creating a Forwarding Default Sequence"

"Creating a Forwarding Installable List" on page 889

Creating a Forwarding Default Sequence

The forwarding default sequence is the default forwarding rule, used for all requests lacking policy instructions. Failover is supported if the sequence (only one is allowed) has more than one member.

Note: The default sequence is completely overridden by policy.

A default forwarding sequence works by allowing healthy hosts to take over for an unhealthy host (one that is failing its DNS resolution or its health check). The sequence specifies the order of failover, with the second host taking over for the first host, the third taking over for the second, and so on).

If all hosts are unhealthy, the operation fails either open or closed, depending upon your settings.

This configuration is generally created and managed through policy. If no forwarding policy applies, you can create a default sequence using the VPM or CPL.

See Also

["Creating a Forwarding Installable List"](#)

Creating a Forwarding Installable List

You can create and install the forwarding installable list using one of the following methods:

- ❑ Text Editor, which allows you to enter the installable list of directives (or copy and paste the contents of an already-created file) directly onto the appliance.
- ❑ A local file, created on your system; the ProxySG can browse to the file and install it.
- ❑ A remote URL, where you placed an already-created file on an FTP or HTTP server to be downloaded to the ProxySG.
- ❑ CLI `inline` command.

When the Forwarding Installable List is installed, it replaces the forwarding configuration on the ProxySG. The configuration remains in effect until overwritten by another installable list; the configuration can be modified or overwritten using CLI commands.

Note: During the time that a forwarding installable list is being compiled and installed, forwarding might not be available. Any transactions that come into the ProxySG during this time might not be forwarded properly.

Installation of forwarding installable lists should be done outside peak traffic times.

To create a forwarding installable list:

1. Select the **Configuration > Forwarding > Forwarding Hosts > Install Forwarding File** tab.
2. From the drop-down list, select the method to use to install the forwarding installable list; click **Install**.

Note: A message is written to the event log when you install a list through the SGOS software.

- **Remote URL:**

Enter the fully-qualified URL, including the filename, where the installable list is located. To view the file before installing it, click **View**. Click **Install**. Examine the installation status that displays; click **OK**.

- **Local File:**

Click **Browse** to display the Local File Browse window. Browse for the installable list file on the local system. Open it and click **Install**. When the installation is complete, a results window opens. View the results, close the window, click **Close**.

- **Text Editor:**

The current configuration is displayed in installable list format. You can customize it or delete it and create your own. Click **Install**. When the installation is complete, a results window opens. View the results, close the window, click **Close**.

Note: The Management Console text editor is a way to enter an installable list for forwarding. It is not a way to enter CLI commands. The directives are understood only by the installable list parser for forwarding.

3. Click **Apply**.

Note: You can create forwarding settings using the CLI `#inline forwarding` command. You can use any of the forwarding directives.

For more information on using inline commands, refer to the *Command Line Interface Reference*.

To delete forwarding settings on the ProxySG:

From the `(config)` prompt, enter the following commands to delete a host, a group, or all hosts and groups from the forwarding configuration:

```
SGOS#(config) forwarding
SGOS#(config forwarding) delete {all | group group_name | host
host_alias}
```

Note: Any host or group in the default sequence (or the WebPulse service configuration) is considered in use by policy. As a result, if you try to delete a host or group while it is in the default sequence or WebPulse service configuration, you will receive an error message. You must remove the host/group from the sequence or service first, then delete.

Chapter 44: Using Policy to Manage Forwarding

After forwarding and the SOCKS gateways are configured, use policy to create and manage forwarding rules. Create forwarding and SOCKS gateway rules in the `<Forward>` layer of the Forwarding Policy file or the VPM Policy file (if you use the VPM).

The separate `<Forward>` layer is provided because the URL can undergo URL rewrites before the request is fetched. This rewritten URL is accessed as a `server_url` and decisions about upstream connections are based on the rewritten URL, requiring a separate layer. All policy commands allowed in the `<Forward>` layer are described below.

Table 44–1 Policy Commands Allowed in the `<Forward>` Layer

Forward	Description
Conditions	
<code>client_address=</code>	Tests the IP address of the client. Can also be used in <code><Exception></code> and <code><Proxy></code> layers.
<code>client.host=</code>	Tests the hostname of the client (obtained through RDNS). Can also be used in <code><Admin></code> , <code><Proxy></code> , and <code><Exception></code> layers.
<code>client.host.has_name=</code>	Tests the status of the RDNS performed to determine <code>client.host</code> . Can also be used in <code><Admin></code> , <code><Proxy></code> , and <code><Exception></code> layers.
<code>client.protocol=</code>	Tests true if the client transport protocol matches the specification. Can also be used in <code><Exception></code> and <code><Proxy></code> layers.
<code>date[.utc]=</code>	Tests true if the current time is within the <code>startdate..enddate</code> range, inclusive. Can be used in all layers.
<code>day=</code>	Tests if the day of the month is in the specified range or an exact match. Can be used in all layers.
<code>has_client=</code>	<code>has_client=</code> is used to test whether or not the current transaction has a client. This can be used to guard triggers that depend on client identity.
<code>hour[.utc]=</code>	Tests if the time of day is in the specified range or an exact match. Can be used in all layers.
<code>im.client=</code>	Tests the type of IM client in use. Can also be used in <code><Proxy></code> , <code><Exception></code> , and <code><Cache></code> layers.
<code>im.message.reflected=</code>	Tests whether IM reflection occurred. Can also be used in <code><Proxy></code> and <code><Cache></code> layers.

Table 44–1 Policy Commands Allowed in the <Forward> Layer (Continued)

Forward	Description
minute[.utc]=month[.utc]=	Tests if the minute of the hour is in the specified range or an exact match. Can be used in all layers.
proxy.address=	Tests the IP address of the network interface card (NIC) on which the request arrives. Can also be used in <Admin> and <Proxy> layers.
proxy.card=	Tests the ordinal number of the network interface card (NIC) used by a request. Can also be used in <Admin> and <Proxy> layers.
proxy.port=	Tests if the IP port used by a request is within the specified range or an exact match. Can also be used in <Admin> and <Proxy> layers.
server_url[.case_sensitive .no_lookup]=	Tests if a portion of the requested URL exactly matches the specified pattern.
server_url.address=	Tests if the host IP address of the requested URL matches the specified IP address, IP subnet, or subnet definition.
server_url.category=	Tests the content categories of the requested URL as assigned by policy definitions or an installed content filter database.
server_url.domain[.case_sensitive .no_lookup]=	Tests if the requested URL, including the domain-suffix portion, matches the specified pattern.
server_url.extension[.case_sensitive]=	Tests if the filename extension at the end of the path matches the specified string.
server_url.host.has_name=	Tests whether the server URL has a resolved DNS hostname.
server_url.host[.exact .substring .prefix .suffix .regex][.no_lookup]=	Tests if the host component of the requested URL matches the IP address or domain name.
server_url.host.is_numeric=	This is true if the URL host was specified as an IP address.
server_url.host.no_name=	This is true if no domain name can be found for the URL host.
server_url.host.regex=	Tests if the specified regular expression matches a substring of the domain name component of the requested URL.
server_url.is_absolute=	Tests whether the server URL is expressed in absolute form.
server_url.path[.exact .substring .prefix .suffix .regex .case_sensitive]=	Tests if a prefix of the complete path component of the requested URL, as well as any query component, matches the specified string.
server_url.path.regex=	Tests if the regex matches a substring of the path component of the request URL.

Table 44–1 Policy Commands Allowed in the <Forward> Layer (Continued)

Forward	Description
<code>server_url.port=</code>	Tests if the port number of the requested URL is within the specified range or an exact match.
<code>server_url.query.regex=</code>	Tests if the regex matches a substring of the query string component of the request URL.
<code>server_url.regex=</code>	Tests if the requested URL matches the specified pattern.
<code>server_url.scheme=</code>	Tests if the scheme of the requested URL matches the specified string.
<code>socks=</code>	This condition is true whenever the session for the current transaction involves SOCKS to the client.
<code>socks.version=</code>	Switches between SOCKS 4/4a and 5. Can also be used in <Exception> and <Proxy> layers.
<code>streaming.client=</code>	<code>yes</code> <code>no</code> . Tests the user agent of a Windows, Real Media, or QuickTime player.
<code>time[.utc]=</code>	Tests if the time of day is in the specified range or an exact match. Can be used in all layers.
<code>tunneled=</code>	<code>yes</code> <code>no</code> . Tests TCP tunneled requests, HTTP CONNECT requests, and unaccelerated SOCKS requests
<code>weekday[.utc]=</code>	Tests if the day of the week is in the specified range or an exact match. Can be used in all layers.
<code>year[.utc]=</code>	Tests if the year is in the specified range or an exact match. Can be used in all layers.
Properties	
<code>access_server()</code>	Determines whether the client can receive streaming content directly from the OCS. Set to <code>no</code> to serve only cached content.
<code>ftp.transport()</code>	Determines the upstream transport mechanism. This setting is not definitive. It depends on the capabilities of the selected forwarding host.
<code>forward()</code>	Determines forwarding behavior. There is a box-wide configuration setting (<code>config>forwarding>failure-mode</code>) for the forward failure mode. The optional specific settings can be used to override the default.
<code>forward.fail_open()</code>	Controls whether the ProxySG appliance terminates or continues to process the request if the specified forwarding host or any designated backup or default cannot be contacted.

Table 44–1 Policy Commands Allowed in the <Forward> Layer (Continued)

Forward	Description
<code>http.refresh.recv.timeout()</code>	Sets the socket timeout for receiving bytes from the upstream host when performing refreshes. Can also be used in <Cache> layers.
<code>http.server.connect_attempts()</code>	Sets the number of attempts to connect performed per-address when connecting to the upstream host.
<code>http.server.recv.timeout()</code>	Sets the socket timeout for receiving bytes from the upstream host. Can also be used in <Proxy> layers.
<code>im.transport()</code>	Sets the type of upstream connection to make for IM traffic.
<code>integrate_new_hosts()</code>	Determines whether to add new host addresses to health checks and load balancing. The default is no. If it is set to yes, any new host addresses encountered during DNS resolution of forwarding hosts are added to health checks and load balancing.
<code>reflect_ip()</code>	Determines how the client IP address is presented to the origin server for explicitly proxied requests. Can also be used in <Proxy> layers.
<code>socks_gateway()</code>	The <code>socks_gateway()</code> property determines the gateway and the behavior of the request if the gateway cannot be contacted. There is a box-wide configuration setting for the SOCKS failure mode. The optional specific settings can be used to override the default.
<code>socks_gateway.fail_open()</code>	Controls whether the ProxySG terminates or continues to process the request if the specified SOCKS gateway or any designated backup or default cannot be contacted.
<code>streaming.transport()</code>	Determines the upstream transport mechanism. This setting is not definitive. The ability to use <code>streaming.transport()</code> depends on the capabilities of the selected forwarding host.
<code>trace.request()</code>	Determines whether detailed trace output is generated for the current request. The default value is no, which produces no output
<code>trace.rules()</code>	Determines whether trace output is generated that shows each policy rule that <i>fired</i> . The default value of no suppresses output.
<code>trace.destination()</code>	Used to change the default path to the trace output file. By default, policy evaluation trace output is written to an object in the cache accessible using a console URL of the following form: <code>http://ProxySG_ip_address:8082/Policy/Trace/path</code>

Table 44–1 Policy Commands Allowed in the <Forward> Layer (Continued)

Forward	Description
Actions	
<code>notify_email()</code>	Sends an e-mail notification to the list of recipients specified in the Event Log mail configuration. Can be used in all layers.
<code>notify_snmp()</code>	The SNMP trap is sent when the transaction terminates. Can be used in all layers.
<code>log_message</code>	Writes the specified string to the event log.
Definitions	
<code>define server_url.domain condition name</code>	Binds a user-defined label to a set of domain suffix patterns for use in a <code>condition= expression</code> .

Chapter 45: About Security

Enterprise-wide security begins with security on the ProxySG, and continues with controlling user access to the Intranet and Internet.

SSH and HTTPS are the recommended (and default) methods for managing access to the ProxySG. SSL is the recommended protocol for communication between the and a realm's off-box authentication server.

Topics in this Section

This section includes information about the following topics:

- ❑ ["Controlling ProxySG Access" on page 897](#)
- ❑ ["Controlling User Access with Identity-based Access Controls" on page 898](#)

Controlling ProxySG Access

You can control access to the ProxySG several ways: by limiting physical access to the system, by using passwords, restricting the use of console account, through per-user RSA public key authentication, and through Blue Coat Content Policy Language (CPL). How secure the system needs to be depends upon the environment.

You can limit access to the ProxySG by:

- ❑ Restricting physical access to the system and by requiring a PIN to access the front panel.
- ❑ Restricting the IP addresses that are permitted to connect to the ProxySG CLI.
- ❑ Requiring a password to secure the Setup Console.

These methods are in addition to the restrictions placed on the console account (a console account user password) and the Enable password. For information on using the console account, see [Chapter 4: "Controlling Access to the ProxySG" on page 59](#) and [Chapter 69: "Configuring Management Services" on page 1267](#).

By using every possible method (physically limiting access, limiting workstation IP addresses, and using passwords), the ProxySG is very secure.

After the ProxySG is secure, you can limit access to the Internet and intranet. It is possible to control access to the network without using authentication. You only need to use authentication if you want to use identity-based access controls.

Controlling User Access with Identity-based Access Controls

The ProxySG provides a flexible authentication architecture that supports multiple services with multiple backend servers (for example, LDAP directory servers together with NT domains with no trust relationship) within each authentication scheme with the introduction of the *realm*.

A *realm* authenticates and authorizes users for access to ProxySG services using either explicit proxy or transparent proxy mode, discussed in "[About Proxy Services](#)" on page 110.

Multiple authentication realms can be used on a single ProxySG. Multiple realms are essential if the enterprise is a managed provider or the company has merged with or acquired another company. Even for companies using only one protocol, multiple realms might be necessary, such as the case of a company using an LDAP server with multiple authentication boundaries. You can use realm sequencing to search the multiple realms all at once.

A realm configuration includes:

- ❑ Realm name.
- ❑ Authentication service—IWA, LDAP, RADIUS, Local, Certificate, Sequences, CA eTrust SiteMinder®, Oracle COREid™, Policy Substitution, Windows SSO, Novell SSO).
- ❑ External server configuration—Backend server configuration information, such as host, port, and other relevant information based on the selected service.
- ❑ Authentication schema—The definition used to authenticate users.
- ❑ Authorization schema—The definition used to authorize users for membership in defined groups and check for attributes that trigger evaluation against any defined policy rules.
- ❑ One-time passwords are supported for RADIUS realms only.

You can view the list of realms already created on the **Configuration > Authentication > Realms** tab. Realms are created on the home page for each realm.

Chapter 46: Controlling Access to the Internet and Intranet

The following sections describe how to limit user access to the Internet and intranet:

- ❑ [Section A: "Managing Users" on page 900](#)
- ❑ [Section B: "Using Authentication and Proxies" on page 907](#)
- ❑ [Section C: "Using SSL with Authentication and Authorization Services" on page 919](#)
- ❑ [Section D: "Creating a Proxy Layer to Manage Proxy Operations" on page 921](#)
- ❑ [Section E: "Forwarding BASIC Credentials" on page 930](#)

Section A: Managing Users

When a user is first authenticated to an ProxySG, a user login is created. You can view users who are logged in and configure the ProxySG to log them out and refresh their data.

This section includes the following topics:

- ❑ ["About User Login"](#) on page 900
- ❑ ["Viewing Logged-In Users"](#) on page 900
- ❑ ["Logging Out Users"](#) on page 901
- ❑ ["Refreshing User Data"](#) on page 903
- ❑ ["Related CLI Syntax to Manage Users"](#) on page 905

About User Login

A user login is the combination of:

- ❑ An IP address
- ❑ A username
- ❑ A realm

For a specific realm, a user is only considered to be logged in once from a given workstation, even if using multiple user agents. However:

- ❑ If policy authenticates the user against multiple realms, the user is logged in once for each realm.
- ❑ If a user logs in from multiple workstations, the user is logged in once per workstation.
- ❑ If multiple users share an IP address (same server, terminal services, or are behind a NAT, which allows a local-area network to use one set of IP addresses), each user is logged in once.
- ❑ If a user logs in from multiple workstations behind a NAT, the user is logged in once.

Viewing Logged-In Users

You can browse all users logged into the ProxySG. You can also filter the displayed users by Glob-username pattern, by IP address subnet, and by realm.

The glob-based username pattern supports three operators:

- ❑ `*` : match zero or more characters
- ❑ `?` : match exactly one character
- ❑ `[x-y]` : match any character in the character range from `x` to `y`

The IP address subnet notation is based on Classless Inter-Domain Routing (CIDR), a way of interpreting IP addresses, as follows:

- ❑ 1.2.3.4: the IP address 1.2.3.4
- ❑ 1.2.3.0/24: the subnet 1.2.3.0 with netmask 255.255.255.0

The realm selection allows an exact realm name or **All realms** to be selected. You can use a combination of these filters to display only the users you are interested in.

To browse users:

1. Select the **Statistics > Authentication** tab.

The screenshot shows a web interface titled "User Logins" with a sub-header "Display User Logins". On the left, three numbered arrows point to the filter fields: arrow 2 points to the "Realm:" dropdown menu which is set to "All realms"; arrow 3 points to the "User pattern:" text input field; and arrow 4 points to the "IP prefix:" text input field. At the bottom of the filter section are two buttons: "Display by user" and "Display by IP".

2. Select a single realm or **All realms** from the **Realm** drop-down list.
3. (Optional) Enter a regular expression in the **User pattern** field to display the usernames that match the pattern.
4. (Optional) Enter an IP address or subnet in the **IP prefix** field to display the IP addresses that match the prefix.
5. Click **Display by user** to display the statistic results by user, or **Display by IP** to display the results by IP address.

Logging Out Users

A logged-in user can be logged out with one of three mechanisms:

- ❑ Inactivity timeout (see "[Inactivity Timeout](#)" on page 902)
- ❑ Explicit logout by the administrator (see "[Administrator Action](#)" on page 902)
- ❑ Policy (see "[Policy](#)" on page 902)

A logged-out user must re-authenticate with the proxy before logging back in.

- ❑ For single sign-on (SSO) realms (Windows SSO, Novell SSO, and IWA configured for SSO), reauthentication is transparent to the user.
- ❑ For non-SSO realms, the user is explicitly challenged for credentials after logout, depending on the **Challenge user after logout** setting in the ProxySG's realm.

Note: The **Challenge user after logout** option only works when cookie-surrogate credentials are used. If this setting is enabled, the user is explicitly challenged for credentials after logging out.

Inactivity Timeout

Each realm has a new inactivity-timeout setting, used in conjunction with the last activity-time value for a particular login. Each time that a login is completed, this activity time is updated. If the time since the last activity time for a specific login exceeds the inactivity-timeout value, the user is logged out.

Administrator Action

The administrator can explicitly log out a set of users using the **Logout** link at the bottom of the user login information pages. See ["Viewing Logged-In Users"](#) on page 900 for information about displaying user login information. For information about using the CLI to logout users, see ["Related CLI Syntax to Manage Users"](#) on page 905.

Policy

Policy has three properties and three conditions to manage user logouts. These properties and conditions can be used to dynamically log out users. For example, you can create a logout link for users.

For information about using policy, refer to the *Visual Policy Manager Reference* and the *Content Policy Language Guide*.

New Properties

Policy has three properties for logging out users.

- ❑ `user.login.log_out (yes)`

This property logs out the user referenced by the current transaction.

- ❑ `user.login.log_out_other (yes)`

If a user is logged in at more than one IP address, this property logs the user out from all IP addresses except the current IP address.

- ❑ `client.address.login.log_out_other (yes)`

If more than one user is logged in at the IP address of the current transaction, this property logs out all users from the current IP address except the current user.

New Conditions

Several conditions support different logout policies.

- ❑ `user.login.count`

This condition matches the number of times that a specific user is logged in with the current realm. You can use this condition to ensure that a user can be logged in only at one workstation. If the condition is combined with the `user.login.log_out_other` property, old login sessions on other workstations are automatically logged out.

- ❑ `client.address.login.count`

This condition matches the number of different users who are logged into the current IP address, and you can use it to limit the user number.

- ❑ `user.login.time`

This condition matches the number of seconds since the current login started, and you can use it to limit the length of a login session.

Refreshing User Data

You can refreshing user data with the following refresh-time options on the specified realm on the ProxySG:

- ❑ Credential refresh time: This option specifies how long a cached username and password is *trusted* (do not require revalidation).
- ❑ Surrogate refresh time: This option specifies how long surrogate credentials are trusted in a particular realm.
- ❑ Authorization refresh time: This option specifies how long authorization data, such as groups and attributes, are trusted.

While the realms have the baseline settings for the different refresh times, policy and administrator actions can override the realm settings. Using the same interface and filters as used for viewing logins, the administrator can select logins and refresh the authorization data, the credentials, or the surrogate credentials using the links available on the user login information page. Refreshing user data might be necessary if users are added to new groups or there is concern about the actual identity of the user on a long-lived IP surrogate credential.

Credential Refresh Time

You can set the credential refresh time with realms that can cache the username and password on the ProxySG. This is limited to realms that use Basic username and password credentials, including LDAP, RADIUS, XML, IWA (with Basic credentials), SiteMinder, and COREid.

Note: The local realm uses Basic credentials but does not need to cache them since they are stored already on the ProxySG.

Cached Usernames and Passwords

You can use a cached username and password to verify a user's credentials without having to verify the credentials with the off-box authentication server. Essentially, this reduces the load on the authentication server. For authentication modes that do not use surrogate credentials (that is, proxy or origin modes), this can greatly reduce the traffic to the authentication server.

The credential refresh time value determines how long a cached username and password is trusted. After that time has expired, the next transaction that needs credential authentication sends a request to the authentication server. A password different than the cached password also results in a request to the authentication server.

One-Time Passwords

One-time passwords are trusted for the credential refresh time. Only when the credential refresh time expires is the user challenged again.

Authorization Refresh Time

Realms (Local, LDAP, Windows SSO, Novell SSO, Certificate, XML, and Policy Substitution) that can do authorization and authentication separately can use the authorization refresh time value to manage the load on the authorization server.

These realms determine authorization data (group membership and attribute values) separately from authentication, allowing the time the authorization data is trusted to be increased or decreased

For realms that must authenticate the user to determine authorization data, the authorization data is updated only when the user credentials are verified by the authentication server.

Surrogate Refresh Time

This value manages how long surrogate credentials are trusted in a particular realm. The authentication mode determines the type of surrogate credential that is used.

- ❑ Cookie surrogate credentials are used with one of the cookie authentication modes; IP address surrogates are used with one of the IP authentications modes; and the Auto authentication mode attempts to select the best surrogate for the current transaction.
- ❑ IP address surrogate credentials work with all user agents, but require that each workstation has a unique IP address; they do not work with users behind a NAT. An IP surrogate credential authenticates all transactions from a given IP address as belonging to the user who was last authenticated at that IP address.

When a user is logged out, all surrogate credentials are discarded, along with the cached credentials and authorization data.

For more information about using cookie and IP address surrogate credentials, see ["About Authentication Modes"](#) on page 908.

Policy

Policy has three properties for setting the refresh times for individual transactions.

❑ `authenticate.authorization_refresh_time(x)`

where *x* is the number of seconds to use for the authorization refresh time during this transaction. The refresh time cannot exceed the time configured in the realm; policy can be used only to reduce the authorization refresh time. You can use this property to dynamically force the user's authorization data to be refreshed.

❑ `authenticate.credential_refresh_time(x)`

where *x* is the number of seconds to use for the credential refresh time during this transaction. The refresh time cannot exceed the time configured in the realm; policy can be used only to reduce the credential refresh time. You can use this property to dynamically force the user's credentials to be refreshed.

❑ `authenticate.surrogate_refresh_time(x)`

where *x* is the number of seconds to use for the surrogate refresh time during this transaction. The refresh time cannot exceed the time configured in the realm; policy can be used only to reduce the surrogate refresh time. You can use this property to dynamically force the user's surrogate to be refreshed.

For information about using policy, refer to the *Visual Policy Manager Reference* and the *Content Policy Language Guide*.

Related CLI Syntax to Manage Users

❑ To enter the manage users submode, use the following commands:

```
SGOS#(config) security users
SGOS#(config users)
```

❑ The following commands are available:

```
(config users) authorization-refresh {ip-addresses prefix [realm_name] | realms [realm_name] | users glob_user_name [realm_name]}
(config users) credentials-refresh {ip-addresses prefix [realm_name] | realms [realm_name] | users glob_user_name [realm_name]}
(config users) log-out {ip-addresses prefix [realm_name] | realms [realm_name] | users glob_user_name [realm_name]}
(config users) surrogates-refresh {ip-addresses prefix [realm_name] | realms [realm_name] | users glob_user_name [realm_name]}
(config users) view {detailed {ip-addresses prefix [realm_name] | realms [realm_name] | users glob_user_name [realm_name]} | ip-addresses prefix [realm_name] | realms [realm_name] | users glob_user_name [realm_name]}
```

Note: Usernames and passwords can each be from 1 to 64 characters in length, but the passwords must be in quotes.

Usernames that contain \ (backward slash), * (asterisk), or ? (question mark) must be escaped when viewing users from the command line interface. The escape character is \.

For example:

- ❑ `user1*` is searched as `#(config users) view users user1\*`
 - ❑ `user1?` is searched as `#(config users) view users user1\?`
 - ❑ `user1\` is searched as `#(config users) view users user1\\`
-

Section B: Using Authentication and Proxies

The ProxySG appliance performs authentication to obtain proof of user identity and then make decisions based on the identity. The appliance obtains proof of identity by sending the client (a browser, for example) a *challenge*—a request to provide credentials. Once the client supplies the credentials, the appliance authenticates (verifies or rejects) them.

Browsers can respond to different kinds of credential challenges:

- ❑ Proxy-style challenges—Sent from proxy servers to clients that are explicitly proxied. In HTTP, the response code is 407.

An authenticating explicit proxy server sends a proxy-style challenge (407/Proxy-Authenticate) to the browser. The browser knows it is talking to a proxy and that the proxy wants proxy credentials. The browser responds to a proxy challenge with proxy credentials (Proxy-Authorization: header). The browser must be configured for explicit proxy in order for it to respond to a proxy challenge.

- ❑ Origin-style challenges—Sent from origin content servers (OCS), or from proxy servers impersonating a OCS. In HTTP, the response code is 401 Unauthorized.

In transparent proxy mode, the ProxySG uses the OCS authentication challenge (HTTP 401 and WWW-Authenticate)—acting as though it is the location from which the user initially requested a page. A transparent proxy, including a reverse proxy, must not use a proxy challenge, because the client might not be expecting it.

- ❑ Client certificate challenges—Sent from servers to initiate an exchange of certificates. In *mutual SSL authentication*, an SSL connection between a client and a server is established only if the client and server validate each other's identity during the SSL handshake. Both parties must have their own valid certificate and the associated private key in order to authenticate.

You might have to configure mutual SSL authentication for an HTTPS reverse proxy service, or the HTTPS-Console service for Common Access Card (CAC) authentication. For information, see "[About Mutual SSL Authentication](#)" on page 325.

Terminology

- ❑ authentication modes: The various ways that the ProxySG interacts with the client for authentication. For more information, see "[About Authentication Modes](#)" on page 908.
- ❑ challenge type: The kind of authentication challenge that is issues (for example, proxy or origin-ip-redirect).
- ❑ guest authentication: Allowing a guest to login with limited permissions.
- ❑ impersonation: The proxy uses the user credentials to connect to another computer and access content that the user is authorized to see.

- ❑ **surrogate credentials:** Credentials accepted in place of the user's real credentials. Surrogate credentials can be either cookie-based or IP address-based.
- ❑ **virtual authentication site:** Used with authentication realms such as IWA, and LDAP. The request for credentials is redirected to the ProxySG instead of the origin server. The appliance intercepts the request for the virtual authentication site and issues the appropriate credential challenge. Thus, the challenge appears to come from the virtual site, which is usually named to make it clear to the user that ProxySG credentials are requested.

About Authentication Modes

Specify an authentication mode to control the way the ProxySG appliance interacts with the client for authentication. The mode specifies the challenge type and the accepted surrogate credential.

- ❑ **Auto:** The default; the mode is automatically selected, based on the request. Auto can choose any of **proxy**, **origin**, **origin-ip**, or **origin-cookie-redirect**, depending on the kind of connection (explicit or transparent) and the transparent authentication cookie configuration.

- ❑ **Proxy:** The ProxySG uses an explicit proxy challenge. No surrogate credentials are used. This is the typical mode for an authenticating explicit proxy. In some situations proxy challenges do not work; origin challenges are then issued.

If you have many requests consulting the back-end authentication authority (such as LDAP, RADIUS, or the BCAAA service), you can configure the ProxySG (and possibly the client) to use persistent connections. This dramatically reduces load on the back-end authentication authority and improves the all-around performance of the network.

- ❑ **Proxy-IP:** The ProxySG uses an explicit proxy challenge and the client's IP address as a surrogate credential. Proxy-IP specifies an insecure forward proxy, possibly suitable for LANs of single-user workstations. In some situations proxy challenges do not work; origin challenges are then issued.
- ❑ **Origin:** The ProxySG acts like an OCS and issues OCS challenges. The authenticated connection serves as the surrogate credential.
- ❑ **Origin-IP:** The ProxySG acts like an OCS and issues OCS challenges. The client IP address is used as a surrogate credential. **Origin-IP** is used to support IWA authentication to the upstream device when the client cannot handle cookie credentials. This mode is primarily used for automatic downgrading, but it can be selected for specific situations.
- ❑ **Origin-cookie:** The ProxySG acts like an origin server and issues origin server challenges. A cookie is used as the surrogate credential. **Origin-cookie** is used in forward proxies to support pass-through authentication more securely than **origin-ip** if the client understands cookies. Only the HTTP and HTTPS protocols support cookies; other protocols are automatically downgraded to **origin-ip**.

This mode could also be used in reverse proxy situations if impersonation (where the proxy uses the user credentials to connect to another computer. and access content that the user is authorized to see).is not possible and the origin server requires authentication.

- ❑ **Origin-cookie-redirect:** The client is redirected to a virtual URL to be authenticated, and cookies are used as the surrogate credential. The ProxySG does not support origin-redirects with the CONNECT method. For forward proxies, only `origin-*-redirect` modes are supported for Kerberos/IWA authentication. (Any other mode uses NTLM authentication.)

Note: During cookie-based authentication, the redirect request to strip the authentication cookie from the URL is logged as a 307 (or 302) TCP_DENIED.

- ❑ **Origin-IP-redirect:** The client is redirected to a virtual URL to be authenticated, and the client IP address is used as a surrogate credential. The ProxySG does not support origin-redirects with the CONNECT method. For forward proxies, only `origin-*-redirect` modes are supported for Kerberos/IWA authentication. (Any other mode uses NTLM authentication.)
- ❑ **SG2:** The mode is selected automatically, based on the request, and uses the SGOS 2.x-defined rules.
- ❑ **Form-IP:** A form is presented to collect the user's credentials. The form is presented whenever the user's credential cache entry expires.
- ❑ **Form-Cookie:** A form is presented to collect the user's credentials. The cookies are set on the OCS domain only, and the user is presented with the form for each new domain. This mode is most useful in reverse proxy scenarios where there are a limited number of domains.
- ❑ **Form-Cookie-Redirect:** A form is presented to collect the user's credentials. The user is redirected to the authentication virtual URL before the form is presented. The authentication cookie is set on both the virtual URL and the OCS domain. The user is only challenged when the credential cache entry expires.
- ❑ **Form-IP-redirect:** This is similar to **form-ip** except that the user is redirected to the authentication virtual URL before the form is presented.

Note: Modes that use an IP address surrogate credential are insecure: After a user has authenticated from an IP address, all further requests from that IP address are treated as from that user. If the client is behind a NAT, or on a multi-user system, this can present a serious security problem.

The default value is `auto`.

For more information about using authentication modes, refer to the *Content Policy Language Guide*.

Setting the Default Authenticate Mode Property

Setting the `authentication.mode` property selects a challenge type and surrogate credential combination. In `auto` mode, explicit IWA uses connection surrogate credentials. In `sg2` mode, explicit IWA uses IP surrogate credentials.

To configure the IWA default authenticate mode settings:

```
SGOS#(config) security default-authenticate-mode {auto | sg2}
```

About Origin-Style Redirection

Some authentication modes redirect the browser to a *virtual authentication site* before issuing the origin-style challenge. This gives the user feedback as to which credentials are required, and makes it possible to (but does not require) send the credentials over a secure connection.

Since browser requests are transparently redirected to the ProxySG, the appliance intercepts the request for the virtual authentication site and issues the appropriate credential challenge. Thus, the challenge appears to come from the virtual site, which is usually named to make it clear to the user that ProxySG credentials are requested.

If authentication is successful, the ProxySG establishes a surrogate credential and redirects the browser back to the original request, possibly with an encoded surrogate credential attached. This allows the ProxySG to see that the request has been authenticated, and so the request proceeds. The response to that request can also carry a surrogate credential.

To provide maximum flexibility, the virtual site is defined by a URL. Requests to that URL (only) are intercepted and cause authentication challenges; other URLs on the same host are treated normally. Thus, the challenge appears to come from a host that in all other respects behaves normally.

Note: Sharing the virtual URL with other content on a real host requires additional configuration if the credential exchange is over SSL.

You can configure the virtual site to something that is meaningful for your company. The default, which requires no configuration, is `www.cfauth.com`. See ["Configuring Transparent Proxy Authentication"](#) on page 911 to set up a virtual URL for transparent proxy.

Tip: Using CONNECT and Origin-Style Redirection

You cannot use the CONNECT method with origin-style redirection or form redirect modes. An error message similar to the following is displayed:

```
Cannot use origin-redirect for CONNECT method (explicit proxy of https URL)
```

Instead, you can add policy to either bypass authentication on the CONNECT method, or use proxy authentication. For example:

```
<proxy>
  allow http.method=CONNECT authenticate.mode(proxy)
authenticate(ldap)
  allow authenticate(cert) authenticate.mode(origin-cookie-redirect)
```

Selecting an Appropriate Surrogate Credential

IP address surrogate credentials are less secure than cookie surrogate credentials and should be avoided if possible. If multiple clients share an IP address (such as when they are behind a NAT firewall or on a multi-user system), the IP surrogate credential mechanism cannot distinguish between those users.

Manually Entering Top-Level Domains (TLDs)

To ensure the proper handling of authentication cookies for top-level domains, a public database (the Public Suffix List) was created. The appliance maintains an internal suffix list for the same purpose. Because there may be instances when the internal list does not properly handle or include new suffixes, a new feature was added to allow administrators to manually add top-level domains.

To manually enter a top-level domain:

1. Select **Configuration > Authentication > Top Level Domains**.
2. Click **Add**.
3. Enter the top-level domain in the Add Domain dialog box.
4. Click **OK**, then **Apply**.

To remove a top-level domain:

1. Select **Configuration > Authentication > Top Level Domains**.
2. Select the top-level domain to remove.
3. Click **Remove**, then **Apply**.

To remove all top-level domains:


1. Select **Configuration > Authentication > Top Level Domains**.
2. Click **Clear All**, then **Apply**.

Configuring Transparent Proxy Authentication

The following sections provide general instructions on configuring for transparent proxy authentication. For more information on transparent proxy, see ["About the Transparent Proxy"](#) on page 105.

To set transparent proxy options:

1. Select the **Configuration > Authentication > Transparent Proxy** tab.



The screenshot shows a configuration window titled "Transparent Proxy". Inside, there is a section labeled "Transparent proxy options". Under the heading "Method:", there are two radio button options: "Cookie" and "IP". The "Cookie" radio button is selected, indicated by a green dot.

2. Select the transparent proxy method—Cookie-based or IP address-based. The default is **Cookie**.
3. Click **Apply**.

Permitting Users to Log in with Authentication or Authorization Failures

You can configure policy (VPM or CPL) to attempt user authentication while permitting specific authentication or authorization errors. The policy can specify that, after certain authentication or authorization failures, the user transaction should be allowed to proceed and not be terminated.

Note: For a list of permitted authentication and authorization errors, see [Chapter 66: "Authentication and Authorization Errors"](#) on page 1211.

Permitted Errors

Authentication and authorization can be permitted to fail if policy has been written to allow specific failures. The behavior is as follows:

- ❑ **Authentication Failures:** After an authentication failure occurs, the authentication error is checked against the list of errors that policy specifies as permitted.
 - If the error is not on the list, the transaction is terminated.
 - If the error is on the list, the transaction is allowed to proceed although the user is unauthenticated. Because the transaction is not considered authenticated, the `authenticated=yes` policy condition evaluates to false and the user has no username, group information, or surrogate credentials. Policy that uses the user, group, domain, or attribute conditions does not match.
- ❑ **Authorization Failures:** After an authorization failure occurs, the authorization error is checked against the list of errors that policy specifies as permitted.
 - If the error is not on the list, the transaction is terminated.
 - If the error is on the list, the transaction is allowed to proceed and the user is marked as not having authorization data.

- If a user is successfully authenticated but does not have authorization data, the `authenticated=yes` condition evaluates to true and the user has valid authentication credentials.
- The `user.authorization_error=any` is evaluate to true if user authorization failed, the user object contains username and domain information, but not group or attribute information. As a result, policy using user or domain actions still match, but policy using group or attribute conditions do not.

To view all authentication and authorization errors, use the `SGOS# show security authentication-errors` CLI command.

Policy Used with Permitted Errors

Before creating policy to permit errors, you must:

- ❑ Identify the type of access the transactions should be permitted.
- ❑ Identify under which circumstances transactions can proceed even if authentication or authorization fails.
- ❑ Identify which errors correspond to those circumstances.

You can use the advanced authentication URL (**Statistics > Advanced > Show Authentication Error Statistics**) as a troubleshooting guide. The policy substitutions `$(x-sc-authentication-error)` and `$(x-sc-authorization-error)` can also be used to log the errors on a per-transaction basis.

Policy conditions and properties that are available include:

- ❑ `authenticate.tolerate_error( )`
- ❑ `authorize.tolerate_error( )`
- ❑ `user.authentication_error=`
- ❑ `user.authorization_error=`
- ❑ `has_authorization_data=`

Note: You are not limited to these conditions and properties in creating policy. For a discussion and a complete list of policy conditions and properties you can use, refer to the *Content Policy Language Guide*.

You can also use the following policy substitutions:

- ❑ `x-sc-authentication-error`: If authentication has failed, this is the error corresponding to the failure. If authentication has not been attempted, the value is `not_attempted`. If authentication has succeeded, the value is `none`.
- ❑ `x-sc-authorization-error`: If authorization has failed, this is the error corresponding to the failure. If authorization has not been attempted, the value is `not_attempted`. If authorization has succeeded, the value is `none`.

Using Guest Authentication

Using policy (VPM or CPL), you can allow a user to log in as a guest user. Guest authentication allows you to assign a username to a user who would have otherwise been considered unauthenticated.

Note: You can use guest authentication with or without default groups. If you use default groups, you can assign guest users to groups for tracking and statistics purposes. For more information about default groups, see ["Using Default Groups"](#) on page 915.

In the case of guest authentication, a user is not actually authenticated against the realm, but is:

- ❑ Assigned the specified guest username
- ❑ Marked as authenticated in the specified realm
- ❑ Marked as a guest user
- ❑ Tracked in access logs

Since the user is not actually authenticated, the username does not have to be valid in that realm.

Using Policy with Guest Authentication

Before creating policy for guest authentication:

- ❑ Determine the circumstances in which guest access is permitted. Guest users are typically allowed in circumstances where no authentication is needed.
- ❑ Determine authentication policy. Will the realms attempt to authenticate users first and fall back to guest authentication, or authenticate users as guest users without attempting authentication?

Note: If a transaction matches both a regular authentication action and guest authentication action, the appliance attempts regular authentication first. This can result in a user challenge before failing over to guest authentication. If a user enters invalid credentials and is thus allowed guest access, they must log out as guest or close and reopen the browser if using session cookies or connection surrogates. They can then enter the correct credentials to obtain regular access.

- ❑ Write the corresponding policy. Policy available for guest authentication includes:
 - `authenticate.guest`
 - `user.is_guest`
 - `authenticated`

Note: You are not limited to these conditions and properties in creating policy. For a complete list of policy conditions and properties you can use, refer to the *Content Policy Language Guide*.

Using Policy Substitutions with Guest Authentication

The following policy substitution was created for use with guest authentication.

- ❑ `x-cs-user-type`: If the user is an authenticated guest user, the value is `guest`. If the user is an authenticated non-guest user, the value is `authenticated`. If the user is not authenticated, the value is `unauthenticated`.

You are not limited to this substitution, and you can use the substitution in other circumstances.

Using Default Groups

You can use default groups with any realm, and they can be used when authorization succeeds, fails or wasn't attempted at all. Default groups allow you to assign users to groups and use those groups in reporting and subsequent authorization decisions.

Note: You can use default groups in conjunction with guest users (see "[Using Guest Authentication](#)" on page 914) or it can be used with regular user authentication.

Using Policy with Default Groups

Before creating policy for default groups, you must determine which set of groups are assigned as default.

You can specify a single or multiple groups here. In most cases, only a single group will be required, but occasionally you might need to assign the user to multiple groups:

- ❑ For extra reporting abilities.
- ❑ If the policy is structured in a way that users should receive the same access as if they belonged in multiple different groups.

Policy available for default groups includes:

- ❑ `group`
- ❑ `authorize.add_group`

Note: You are not limited to these conditions and properties in creating policy. For a complete list of policy conditions and properties you can use, refer to the *Content Policy Language Guide*.

Guest Authentication Example

In this scenario, the administrator has already created a realm against which to authenticate users. The administrator wants the ProxySG appliance to authenticate all users, and display an error if authentication fails. If authentication fails, users can log in as guests.

The following example provides instructions that the administrator could use to create policy in the VPM.

- ❑ ["Step 1 - Create a Web Authentication Policy Layer"](#) on page 916
- ❑ ["Step 2 - Authenticate Users"](#) on page 916
- ❑ ["Step 3 - Specify Permitted Authentication Errors"](#) on page 917
- ❑ ["Step 4 - Authenticate Guests"](#) on page 918
- ❑ ["Step 5 - Restrict Guests' Access"](#) on page 918

Step 1 - Create a Web Authentication Policy Layer

Create the Web Authentication Layer and add a Combined Action Object.

1. In the Management Console, select **Configuration > Policy > Visual Policy Manager**.
2. Click **Launch**. The VPM opens.
3. Add a new Web Authentication Layer. Select **Policy > Add Web Authentication Layer**. Name the layer and click **OK**.
4. Right click in the **Action** column. The VPM displays a menu.
5. On the menu, click **Set**. The VPM displays a Set Action Object dialog.
6. Create policy to authenticate users. Select **New > Combined Action Object**. The VPM displays an Add Combined Action Object dialog.

Step 2 - Authenticate Users

Specify how to authenticate users. This is the first action in the Combined Action Object.

1. On the Add Combined Action Object dialog, create policy to authenticate users.

Select **New > Authenticate**. The VPM displays an Add Authenticate Object dialog.

2. Specify the following for the Authenticate object:

- **Name:** A name for the object.
- **Realm:** The realm against which users will authenticate.
- **Mode:** The authentication mode; see ["About Authentication Modes"](#) on page 908 for descriptions of the modes.

Selecting a Form for the authentication mode enables the following settings:

- **Authentication Form:** The form used to challenge users.

- **New PIN Form:** (RSA SecurID only) The form used to prompt users to enter a PIN.
- **Query Form:** (RSA SecurID only) The form used to display a Yes/No question to users.

For more information on these settings, refer to the “The Visual Policy Manager” chapter in the *Visual Policy Manager Reference*.

Click **OK** to save the settings. The dialog lists the action.

3. Select the action and click **Add**. The Selected Action Objects section displays the action.

Step 3 - Specify Permitted Authentication Errors

Specify what happens when a user fails to authenticate. This is the second action in the Combined Action Object create policy to authenticate users.

1. On the Add Combined Action Object dialog, create policy for permitted authentication errors.

Select **New > Permit Authentication Error**. The VPM displays an Add Permit Authentication Error Object dialog.

2. Specify which authentication errors are allowed.

You can select an option beside **Show** to display more or fewer error types.

- **Any error:** Allow all authentication errors. Any authentication error results in failover to guest authentication, and users can log in as guest.
- **Selected errors:** Allow only the specified types of authentication errors. Only the specified authentication errors fail over to guest authentication. All other authentication errors result in the request being denied.

Note: Symantec recommends that you select **Show > All errors** to display all errors. Then, determine exactly which errors to allow and select only those using the **Selected errors** option. Do not select **Any**, because doing so means that the `need_credentials` error (in the **User Credentials Required** group) would be permitted and the appliance would not challenge users. This could result in all users being authenticated as guests, even domain users.

For more information on these settings, refer to the “The Visual Policy Manager” chapter in the *Visual Policy Manager Reference*.

Click **OK** to save the settings. The dialog lists the dialog.

3. Select the object and click **Add**. The Selected Action Objects section displays the action. Make sure that the object is below the Authenticate object.
4. Click **OK**. The Set Action Objects dialog displays the Combined Action Object.
5. Select the object and click **OK**.

6. Click **Install Policy**.

Step 4 - Authenticate Guests

Allow non-authenticated users to log in as guests.

1. Add another Web Authentication Layer. Select **Policy > Add Web Authentication Layer**. Name the layer and click **OK**.
2. Right click in the **Action** column. The VPM displays a menu.
3. On the menu, click **Set**. The VPM displays a Set Action Object dialog.
4. Create policy to authenticate guests. Select **New > Authenticate Guest**. The VPM displays an Add Authenticate Guest Object dialog.
5. Specify the following for the Authenticate Guest object.
 - **Name:** A name for the object.

Guest Username: A name designated for guests; access logs display this name. Accept the defaults for the remaining settings. For more information, settings, refer to the “The Visual Policy Manager” chapter in the *Visual Policy Manager Reference*.

Click **OK** to save the settings. The dialog displays the object.

Step 5 - Restrict Guests' Access

Determine which transactions guests can perform, and then create policy to control guests' access.

1. Create a Web Access Layer. On the VPM dialog, select **Policy > Add Web Access Layer**. Name the layer and click **OK**.
2. Create a condition for guest users by creating a layer guard. Right click the layer name.
3. On the menu that displays, click **Add Layer Guard**. The first row of the layer displays the layer guard.
4. In the layer guard, right click **Source**. The VPM displays the Set Source Object dialog.
5. Select **Guest User** and click **OK**. When policy matches the condition (the source is a guest), the appliance evaluates all of the rules in the layer in the transaction.
6. Add rules to the layer to limit guests' access.
7. Install the policy. On the VPM dialog, click **Install Policy**. The VPM indicates that policy was installed.

Section C: Using SSL with Authentication and Authorization Services

Blue Coat recommends that you use SSL during authentication to secure your user credentials. Blue Coat supports SSL between the client and the ProxySG and between the ProxySG and LDAP and IWA authentication servers as described in the following sections:

- ❑ ["Using SSL Between the Client and the ProxySG"](#) on page 919
- ❑ ["Using SSL Between the ProxySG and the Authentication Server"](#) on page 919

Using SSL Between the Client and the ProxySG

To configure SSL between the client and the ProxySG using `origin-cookie-redirect` or `origin-ip-redirect` challenges, you must:

- ❑ Specify a virtual URL with the HTTPS protocol (for example, `https://virtual_address`).
- ❑ Create a keyring and certificate on the ProxySG.
- ❑ Create an HTTPS service to run on the port specified in the virtual URL and to use the keyring you just created.

Note: You can use SSL between the client and the ProxySG for origin-style challenges on transparent and explicit connections (SSL for explicit proxy authentication is not supported).

In addition, if you use a forward proxy, the challenge type must use redirection; it cannot be an `origin` or `origin-ip` challenge type.

When redirected to the virtual URL, the user is prompted to accept the certificate offered by the ProxySG (unless the certificate is signed by a trusted certificate authority). If accepted, the authentication conversation between the ProxySG and the user is encrypted using the certificate.

Note: If the hostname does not resolve to the IP address of the ProxySG, then the network configuration must redirect traffic for that port to the appliance. Also, if you use the IP address as the virtual hostname, you might have trouble getting a certificate signed by a CA-Certificate authority (which might not be important).

Using SSL Between the ProxySG and the Authentication Server

SSL communication between the ProxySG and LDAP and IWA authentication servers is supported. You configure the SSL communication settings when configuring the realm:

- ❑ For information on configuring SSL between the ProxySG and an IWA realm, see ["Configuring IWA Servers"](#) on page 1020.
- ❑ For information on configuring SSL between the ProxySG and an LDAP realm, see ["Configuring LDAP Servers"](#) on page 1048.

Note: If the browser is configured for online checking of certificate revocation, the status check must be configured to bypass authentication.

Section D: Creating a Proxy Layer to Manage Proxy Operations

Once hardware configuration is complete and the system configured to use transparent or explicit proxies, use CPL or VPM to provide on-going management of proxy operations.

Using CPL

Below is a table of all commands available for use in proxy layers of a policy. If a condition, property, or action does not specify otherwise, it can be used only in <Proxy> layers. For information about creating effective CPL, refer to the *Content Policy Language Guide*.

Table 46–1 CPL Commands Available in the <Proxy> Layer

<Proxy> Layer Conditions	Meaning
admin.access=	Tests the administrative access requested by the current transaction. Can also be used in <Admin> layers.
attribute.name=	Tests if the current transaction is authenticated in a RADIUS or LDAP realm, and if the authenticated user has the specified attribute with the specified value. Can also be used in <Admin> layers.
authenticated=	Tests if authentication was requested and the credentials could be verified; otherwise, false. Can also be used in <Admin> layers.
bitrate=	Tests if a streaming transaction requests bandwidth within the specified range or an exact match. Can also be used in <Cache> layers.
category=	Tests if the content categories of the requested URL match the specified category, or if the URL has not been categorized. Can also be used in <Cache> layers.
client_address=	Tests the IP address of the client. Can also be used in <Admin> layers.
client.connection.negotiated_cipher=	Test the cipher suite negotiated with a securely connected client. Can also be used in <Exception> layers.
client.connection.negotiated_cipher.strength=	Test the cipher strength negotiated with a securely connected client. Can also be used in <Exception> layers.
client.host=	Test the hostname of the client (obtained through RDNS). Can also be used in <Admin>, <Forward>, and <Exception> layers.
client.host.has_name=	Test the status of the RDNS performed to determine client.host. Can also be used in <Admin>, <Forward>, and <Exception> layers.
client_protocol=	Tests true if the client transport protocol matches the specification. Can also be used in <Exception> layers.
condition=	Tests if the specified defined condition is true. Can be used in all layers.

Table 46–1 CPL Commands Available in the <Proxy> Layer (Continued)

<code>console_access=</code>	(This trigger was formerly <code>admin=yes no</code> .) Tests if the current request is destined for the admin layer. Can also be used in <Cache> and <Exception> layers.
<code>content_management=</code>	(This trigger was formerly <code>content_admin=yes no</code> .) Tests if the current request is a content-management transaction. Can also be used in <Exception> and <Forward> layers.
<code>date[.utc]=</code>	Tests true if the current time is within the <code>startdate..enddate</code> range, inclusive. Can be used in all layers.
<code>day=</code>	Tests if the day of the month is in the specified range or an exact match. Can be used in all layers.
<code>exception.id=</code>	Indicates that the requested object was not served, providing this specific exception page. Can also be used in <Exception> layers.
<code>ftp.method=</code>	Tests ftp request methods against any of a well-known set of FTP methods. Can also be used in <Cache> and <Exception> layers.
<code>group=</code>	Tests if the authenticated condition is set to yes, the client is authenticated, and the client belongs to the specified group. Can also be used in <Admin> layers.
<code>has_attribute.name=</code>	Tests if the current transaction is authenticated in an LDAP realm and if the authenticated user has the specified LDAP attribute. Can also be used in <Admin> layers.
<code>hour=</code>	Tests if the time of day is in the specified range or an exact match. Can be used in all layers.
<code>http.method=</code>	Tests HTTP request methods against any of a well known set of HTTP methods. Can also be used in <Cache> and <Exception> layers.
<code>http.method.regex=</code>	Test the HTTP method using a regular expression. Can also be used in <Exception> layers.
<code>http.request_line.regex=</code>	Test the HTTP protocol request line. Can also be used in <Exception> layers.
<code>http.request.version=</code>	Tests the version of HTTP used by the client in making the request to the ProxySG. Can also be used in <Cache> and <Exception> layers.
<code>http.response_code=</code>	Tests true if the current transaction is an HTTP transaction and the response code received from the origin server is as specified. Can also be used in <Cache> and <Exception> layers.
<code>http.response.version=</code>	Tests the version of HTTP used by the origin server to deliver the response to the ProxySG. Can also be used in <Cache> and <Exception> layers.
<code>http.transparent_authentication=</code>	This trigger evaluates to true if HTTP uses transparent proxy authentication for this request. Can also be used in <Cache> and <Exception> layers.

Table 46–1 CPL Commands Available in the <Proxy> Layer (Continued)

<code>im.buddy_id=</code>	Tests the <code>buddy_id</code> associated with the IM transaction. Can also be used in <Exception> layers.
<code>im.chat_room.conference=</code>	Tests whether the chat room associated with the transaction has the <code>conference</code> attribute set. Can also be used in <Exception> layers.
<code>im.chat_room.id=</code>	Tests the chat room ID associated with the transaction. Can also be used in <Exception> layers.
<code>im.chat_room.invite_only=</code>	Tests whether the chat room associated with the transaction has the <code>invite_only</code> attribute set. Can also be used in <Exception> layers.
<code>im.chat_room.type=</code>	Tests whether the chat room associated with the transaction is public or private. Can also be used in <Exception> layers.
<code>im.chat_room.member=</code>	Tests whether the chat room associated with the transaction has a member matching the specified criterion. Can also be used in <Exception> layers.
<code>im.chat_room.voice_enabled=</code>	Tests whether the chat room associated with the transaction is voice enabled. Can also be used in <Exception> layers.
<code>im.client=</code>	Test the type of IM client in use. Can also be used in <Exception>, <Forward>, and <Cache> layers.
<code>im.file.extension=</code>	Tests the file extension. Can also be used in <Exception> layers.
<code>im.file.name=</code>	Tests the file name (the last component of the path), including the extension. Can also be used in <Exception> layers.
<code>im.file.path=</code>	Tests the file path against the specified criterion. Can also be used in <Exception> layers.
<code>im.file.size=</code>	Performs a signed 64-bit range test. Can also be used in <Exception> layers.
<code>im.message.reflected</code>	Test whether IM reflection occurred. Can also be used in <Exception> and <Forward> layers.
<code>im.message.route=</code>	Tests how the IM message reaches its recipients. Can also be used in <Exception> layers.
<code>im.message.size=</code>	Performs a signed 64-bit range test. Can also be used in <Exception> layers.
<code>im.message.text.substring=</code>	Tests if the message text contains the specified text or pattern. Can also be used in <Exception> layers.
<code>im.message.opcode=</code>	Tests the value of an opcode associated with an <code>im.method</code> of <code>unknown_send</code> or <code>unknown_receive</code> .
<code>im.message.type=</code>	Tests the message type. Can also be used in <Exception> layers.
<code>im.method=</code>	Tests the method associated with the IM transaction. Can also be used in <Cache> and <Exception> layers.

Table 46–1 CPL Commands Available in the <Proxy> Layer (Continued)

<code>im.user_id=</code>	Tests the <code>user_id</code> associated with the IM transaction. Can also be used in <Exception> layers.
<code>live=</code>	Tests if the streaming content is a live stream. Can also be used in <Cache> layers.
<code>minute=</code>	Tests if the minute of the hour is in the specified range or an exact match. Can be used in all layers.
<code>month=</code>	Tests if the month is in the specified range or an exact match. Can be used in all layers.
<code>proxy.address=</code>	Tests the IP address of the network interface card (NIC) on which the request arrives. Can also be used in <Admin> layers.
<code>proxy.card=</code>	Tests the ordinal number of the network interface card (NIC) used by a request. Can also be used in <Admin> layers.
<code>proxy.port=</code>	Tests if the IP port used by a request is within the specified range or an exact match. Can also be used in <Admin> layers.
<code>raw_url</code>	Test the value of the raw request URL. Can also be used in <Exception> layers.
<code>raw_url.host</code>	Test the value of the 'host' component of the raw request URL. Can also be used in <Exception> layers.
<code>raw_url.path</code>	Test the value of the 'path' component of the raw request URL. Can also be used in <Exception> layers.
<code>raw_url.pathquery</code>	Test the value of the 'path and query' component of the raw request URL. Can also be used in <Exception> layers.
<code>raw_url.port</code>	Test the value of the 'port' component of the raw request URL. Can also be used in <Exception> layers.
<code>raw_url.query</code>	Test the value of the 'query' component of the raw request URL. Can also be used in <Exception> layers.
<code>realm=</code>	Tests if the authenticated condition is set to yes, the client is authenticated, and the client has logged into the specified realm. Can also be used in <Admin> layers.
<code>release.id=</code>	Tests the ProxySG release ID. Can be used in all layers.
<code>request.header.address. header_name=</code>	Tests if the specified request header can be parsed as an IP address. Can also be used in <Cache> layers.
<code>request.header.header_ name=</code>	Tests the specified request header (<code>header_name</code>) against a regular expression. Can also be used in <Cache> layers.
<code>request.header.header_ name.count</code>	Test the number of header values in the request for the given <code>header_name</code> . Can also be used in <Exception> layers.
<code>request.header.header_ name.length</code>	Test the total length of the header values for the given <code>header_name</code> . Can also be used in <Exception> layers.

Table 46–1 CPL Commands Available in the <Proxy> Layer (Continued)

<code>request.header.Referer.url.host.has_name=</code>	Test whether the Referer URL has a resolved DNS hostname. Can also be used in <Exception> layers.
<code>request.header.Referer.url.is_absolute</code>	Test whether the Referer URL is expressed in absolute form. Can also be used in <Exception> layers.
<code>request.raw_headers.count</code>	Test the total number of HTTP request headers. Can also be used in <Exception> layers.
<code>request.raw_headers.length</code>	Test the total length of all HTTP request headers. Can also be used in <Exception> layers.
<code>request.raw_headers.regex</code>	Test the value of all HTTP request headers with a regular expression. Can also be used in <Exception> layers.
<code>request.x_header.header_name.count</code>	Test the number of header values in the request for the given <i>header_name</i> . Can also be used in <Exception> layers.
<code>request.x_header.header_name.length</code>	Test the total length of the header values for the given <i>header_name</i> . Can also be used in <Exception> layers.
<code>response.header.header_name=</code>	Tests the specified response header (<i>header_name</i>) against a regular expression. Can also be used in <Cache> layers.
<code>response.x_header.header_name=</code>	Tests the specified response header (<i>header_name</i>) against a regular expression. Can also be used in <Cache> layers.
<code>server_url[.case_sensitive .no_lookup]=</code>	Tests if a portion of the requested URL exactly matches the specified pattern. Can also be used in <Forward> layers.
<code>socks.accelerated=</code>	Controls the SOCKS proxy handoff to other protocol agents.
<code>socks.method=</code>	Tests the protocol method name associated with the transaction. Can also be used in <Cache> and <Exception> layers.
<code>socks.version=</code>	Switches between SOCKS 4/4a and 5. Can also be used in <Exception> and <Forward> layers.
<code>streaming.content=</code>	(This trigger has been renamed from streaming.) Can also be used in <Cache>, <Exception>, and <Forward> layers.
<code>time=</code>	Tests if the time of day is in the specified range or an exact match. Can be used in all layers.
<code>tunneled=</code>	
<code>url.domain=</code>	Tests if the requested URL, including the domain-suffix portion, matches the specified pattern. Can also be used in <Forward> layers.
<code>url.extension=</code>	Tests if the filename extension at the end of the path matches the specified string. Can also be used in <Forward> layers.
<code>url.host=</code>	Tests if the host component of the requested URL matches the IP address or domain name. Can also be used in <Forward> layers.

Table 46–1 CPL Commands Available in the <Proxy> Layer (Continued)

<code>url.host.has_name</code>	Test whether the request URL has a resolved DNS hostname. Can also be used in <Exception> layers
<code>url.is_absolute</code>	Test whether the request URL is expressed in absolute form. Can also be used in <Exception> layers
<code>url.host.is_numeric=</code>	This is true if the URL host was specified as an IP address. Can also be used in <Forward> layers.
<code>url.host.no_name=</code>	This is true if no domain name can be found for the URL host. Can also be used in <Forward> layers.
<code>url.host.regex=</code>	Tests if the specified regular expression matches a substring of the domain name component of the request URL. Can also be used in <Forward> layers.
<code>url.host.suffix=</code>	Can also be used in <Forward> layers.
<code>url.path=</code>	Tests if a prefix of the complete path component of the requested URL, as well as any query component, matches the specified string. Can also be used in <Forward> layers.
<code>url.path.regex=</code>	Tests if the regex matches a substring of the path component of the request URL. Can also be used in <Forward> layers.
<code>url.port=</code>	Tests if the port number of the requested URL is within the specified range or an exact match. Can also be used in <Forward> layers.
<code>url.query.regex=</code>	Tests if the regex matches a substring of the query string component of the request URL. Can also be used in <Forward> layers.
<code>url.regex=</code>	Tests if the requested URL matches the specified pattern. Can also be used in <Forward> layers.
<code>url.scheme=</code>	Tests if the scheme of the requested URL matches the specified string. Can also be used in <Forward> layers.
<code>user=</code>	Tests the authenticated user name of the transaction. Can also be used in <Admin> layers.
<code>user.domain=</code>	Tests if the authenticated condition is set to yes, the client is authenticated, the logged-into realm is an IWA realm, and the domain component of the user name is the specified domain. Can also be used in <Admin> layers.
<code>weekday=</code>	Tests if the day of the week is in the specified range or an exact match. Can be used in all layers.
<code>year=</code>	Tests if the year is in the specified range or an exact match. Can be used in all layers.

Table 46–2 Properties Available in the <Proxy> Layer

<Proxy> Layer Properties	Meaning
<code>action.action_label()</code>	Selectively enables or disables a specified define action block. Can also be used in <Cache> layers.
<code>allow</code>	Allows the transaction to be served. Can be used in all layers except <Exception> and <Forward> layers.
<code>always_verify()</code>	Determines whether each request for the objects at a particular URL must be verified with the origin server.
<code>authenticate()</code>	Identifies a realm that must be authenticated against. Can also be used in <Admin> layers.
<code>authenticate.force()</code>	Either disables proxy authentication for the current transaction (using the value <code>no</code>) or requests proxy authentication using the specified authentication realm. Can also be used in <Admin> layers.
<code>authenticate.form()</code>	When forms-based authentication is in use, <code>authenticate.form()</code> selects the form used to challenge the user.
<code>authenticate.mode(auto)</code> <code>authenticate.mode(sg2)</code>	Setting the <code>authenticate.mode</code> property selects a challenge type and surrogate credential combination. In <code>auto</code> mode, explicit IWA uses connection surrogate credentials. In <code>sg2.mode</code> , explicit IWA uses IP surrogate credentials.
<code>authenticate.redirect_stored_requests</code>	Sets whether requests stored during forms-based authentication can be redirected if the upstream host issues a redirecting response.
<code>bypass_cache()</code>	Determines whether the cache is bypassed for a request.
<code>check_authorization()</code>	In connection with CAD (Caching Authenticated Data) and CPAD (Caching Proxy Authenticated Data) support, <code>check_authorization()</code> is used when you know that the upstream device will sometimes (not always or never) require the user to authenticate and be authorized for this object. Can also be used in <Cache> layers.
<code>delete_on_abandonment()</code>	If set to yes, then if all clients requesting an object close their connections prior to the object being delivered, the object fetch from the origin server is abandoned. Can also be used in <Cache> layers.
<code>deny</code>	Denies service. Can be used in all layers except <Exception> and <Forward> layers.
<code>dynamic_bypass()</code>	Used to indicate that a particular transparent request should not be handled by the proxy, but instead be subjected to our dynamic bypass methodology.
<code>exception()</code>	Indicates not to serve the requested object, but instead serve this specific exception page. Can be used in all layers except <Exception> layers.
<code>ftp.server_connection()</code>	Determines when the control connection to the server is established.

Table 46-2 Properties Available in the <Proxy> Layer (Continued)

<code>ftp.welcome_banner()</code>	Sets the welcome banner for a proxied FTP transaction.
<code>http.client.recv.timeout</code>	Sets the socket timeout for receiving bytes from the client.
<code>http.request.version()</code>	The <code>http.request.version()</code> property sets the version of the HTTP protocol to be used in the request to the origin content server or upstream proxy. Can also be used in <Cache> layers.
<code>http.response.parse_meta_tag. Cache-Control()</code>	Controls whether the <code>Cache-Control</code> META Tag is parsed in an HTML response body. Can also be used in <Cache> layers.
<code>http.response.parse_meta_tag. Expires</code>	Controls whether the <code>Expires</code> META Tag is parsed in an HTML response body. Can also be used in <Cache> layers.
<code>http.response.parse_meta_tag. Pragma.no-cache</code>	Controls whether the <code>Pragma: no-cache</code> META Tag is parsed in an HTML response body. Can also be used in <Cache> layers.
<code>http.response.version()</code>	The <code>http.response.version()</code> property sets the version of the HTTP protocol to be used in the response to the client's user agent.
<code>http.server.recv. timeout()</code>	Sets the socket timeout for receiving bytes from the upstream host. Can also be used in <Forward> layers.
<code>im.block_encryption</code>	Prevents the encryption of AOL IM messages by modifying messages during IM login time.
<code>im.reflect</code>	Sets whether IM reflection should be attempted.
<code>im.strip_attachments()</code>	Determines whether attachments are stripped from IM messages.
<code>im.transport</code>	Sets the type of upstream connection to make for IM traffic.
<code>log.suppress.field-id()</code>	The <code>log.suppress.field-id()</code> controls suppression of the specified field-id in all facilities (individual logs that contain all properties for that specific log in one format). Can be used in all layers.
<code>log.suppress.field-id [log_list]()</code>	The <code>log.suppress.field-id [log_list]()</code> property controls suppression of the specified field-id in the specified facilities. Can be used in all layers.
<code>log.rewrite.field-id()</code>	The <code>log.rewrite.field-id()</code> property controls rewrites of a specific log field in all facilities. Can be used in all layers.
<code>log.rewrite.field-id [log_list]()</code>	The <code>log.rewrite.field-id [log_list]()</code> property controls rewrites of a specific log field in a specified list of log facilities. Can be used in all layers.
<code>reflect_ip()</code>	Determines how the client IP address is presented to the origin server for explicitly proxied requests. Can also be used in <Forward> layers.
<code>request.icap_service()</code>	Determines whether a request from a client should be processed by an external ICAP service before going out.
<code>shell.prompt</code>	Sets the prompt for a proxied Shell transaction.
<code>shell.realm_banner</code>	Sets the realm banner for a proxied Shell transaction.

Table 46–2 Properties Available in the <Proxy> Layer (Continued)

<code>shell.welcome_banner</code>	Sets the welcome banner for a proxied Shell transaction.
<code>socks.accelerate()</code>	The <code>socks.accelerate</code> property controls the SOCKS proxy handoff to other protocol agents.
<code>socks.authenticate()</code>	The same realms can be used for SOCKS proxy authentication as can be used for regular proxy authentication.
<code>socks.authenticate.force()</code>	The <code>socks.authenticate.force()</code> property forces the realm to be authenticated through SOCKS.

Table 46–3 Actions Available in the <Proxy> Layer

<Proxy> Layer Actions	Meaning
<code>log_message()</code>	Writes the specified string to the ProxySG event log. Can be used in all layers except <Admin>.
<code>notify_email()</code>	Sends an e-mail notification to the list of recipients specified in the Event Log mail configuration. Can be used in all layers.
<code>notify_snmp()</code>	The SNMP trap is sent when the transaction terminates. Can be used in all layers.
<code>redirect()</code>	Ends the current HTTP transaction and returns an HTTP redirect response to the client.
<code>transform</code>	Invokes the active content or URL rewrite transformer.

Section E: Forwarding BASIC Credentials

Forwarding BASIC credentials enables single sign on when other, more secure, options are unavailable.

About Forwarding BASIC Credentials

Depending upon the application and its configuration, an OCS might require BASIC credentials in order to authenticate the connection from the ProxySG. These credentials can be a fixed username and password or the same credentials used for proxy authentication. Using policy, you can forward these user credentials or send fixed credentials to authenticate the user to the OCS.

Setting Up Policy to Forward Basic Credentials

The policy procedure below assumes no existing policy layers. A properly set up Visual Policy Manager has many existing layers and policies with a logical order. For existing deployments, it is necessary to add new actions to existing layers to enable forwarding BASIC credentials. Make sure you have thoroughly read and are familiar with creating policies before continuing.


Note: Refer to the *Visual Policy Manager Reference* for complete details about the VPM.

Situation


An internal reverse proxy setup. The administrator wishes to forward BASIC credential, either user or custom credentials to a particular OCS.

To forward BASIC credentials:


1. Select the **Configuration > Policy > Visual Policy Manager** tab.
2. Click **Launch**. The VPM launches in a separate window.




3. Select **Policy > Add Web Authentication Layer**. An **Add New Layer** dialog box displays.
4. Enter a name that is easily recognizable and click **OK**. A new policy tab and rule displays in the workspace.



5. Select **None** under the **Action** column. Right click **Any** > **Set**. The **Set Action Object** window displays.



6. Select **New** > **Send Credentials Upstream**.



7. The **Add Send Credentials Upstream Object** window allows configuration of forwarding BASIC credentials.
 - a. Enter an easily recognizable name in the field.
 - b. Select the authentication method from the **Authentication Type** drop-down list. Select **origin** or **proxy**. If you are authenticating to an upstream origin server, select **origin**. If you are authenticating to a proxy server, select **proxy**.

- c. Select the credentials required for a particular OCS.
 - Select the **Send user credentials** radio button to send user credentials to the OCS.
 - Select the **Send custom credentials** radio button to forward a fixed username and password to an OCS. Selection of this option requires the **UserName** and **Password** fields to be filled with the appropriate values.
8. Click **OK**.
9. Click **OK** to return to the VPM.
10. Click the **Install Policy** button when finished adding policies.

Note: For all transactions which match the **Send Credentials Upstream Object**, credentials will be sent even if the receiving server does not require them. Depending upon how your policy is written, you can use the **Do Not Send Credentials Upstream** object to manage which servers should not receive credentials. You can enforce this rule using the VPM object, **Do Not Send Credentials Upstream**. It is a fixed action and requires no configuration.

Creating the CPL

Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file layers.

- ❑ Authenticate to an upstream server using the user's BASIC credentials.

```
<proxy>
url.host.exact="webmail.company.com" server.authenticate.basic(origin)
```
- ❑ Authenticate to an upstream proxy using a fixed username and password.

```
<proxy>
url.host.exact="proxy.company.com" \
server.authenticate.basic(proxy,"internaluser", "internalpassword")
```
- ❑ Authenticate to an upstream server using the IP address of the client.

```
<proxy>
url.host.exact="images.company.com" \
server.authenticate.basic(origin,"$(client.address)")
```

Chapter 47: Local Realm Authentication and Authorization

Using a Local realm is appropriate when the network topology does not include external authentication or when you want to add users and administrators to be used by the ProxySG only.

The Local realm (you can create up to 40) uses a *Local User List*, a collection of users and groups stored locally. You can create up to 50 different Local User Lists. Multiple Local realms can reference the same list at the same time, although each realm can only reference one list at a time. The default list used by the realm can be changed at any time.

Local realm authentication can be used to authenticate administrative users to the ProxySG appliance management console, and is highly recommended. Because the user details are stored on the appliance, local authentication realms are always available.

Topics in this Section


This section includes information about the following topics:

- ❑ ["Creating a Local Realm"](#)
- ❑ ["Changing Local Realm Properties"](#) on page 934
- ❑ ["Defining the Local User List"](#) on page 936
- ❑ ["Creating the CPL"](#) on page 942

Creating a Local Realm

To create a local realm:

1. Select the **Configuration > Authentication > Local > Local Realms** tab.
2. Click **New**.



3. Create the realm:
 - a. In the **Realm name** field, enter a realm name. The name can be 32 characters long and composed of alphanumeric characters and underscores. The name must start with a letter.
 - b. Click **OK**.
4. Click **Apply**.

Changing Local Realm Properties

Once you have created a Local realm, you can modify the properties.

To define or change local realm properties:

1. Select the **Configuration > Authentication > Local > Local Main** tab.

The screenshot shows the 'Local Main' configuration page for a realm named 'Local_test22'. The settings are as follows:

- 2a** Realm name: Local_test22
- 2b** Display name: Local_test22
- 2c** Local user list: local_user_database
- 3a** Refresh Times: ☒ Use the same refresh time for all
- 3b** Credential refresh time: 900 seconds
- 3c** Surrogate refresh time: 900 seconds
- 3d** Authorization refresh time: 900 seconds
- 4** Inactivity timeout: 900 seconds
- 5** Cookies:
 - ☐ Use persistent cookies
 - ☒ Verify the IP address in the cookie
- 6** Virtual URL: www.cfauth.com/
- 7** ☒ Challenge user after logout

2. Configure basic information:

- From the **Realm name** drop-down list, select the **Local** realm for which you want to change properties.
- Display name:** The default value for the display name is the realm name. The display name cannot be greater than 128 characters and it cannot be null.
- Local user list:** the local user list from the drop-down list.

3. Configure refresh options:

- Select the **Use the same refresh time for all** check box if you would like to use the same refresh time for all.
- Enter the number of seconds in the **Surrogate refresh time** field. The Surrogate Refresh Time allows you to set a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the ProxySG authenticates the transaction. After the refresh time expires, the ProxySG verifies the user's credentials. Depending upon the authentication mode and the user-agent, this may result in challenging the end user for credentials.

The main goal of this feature is to verify that the user-agent still has the appropriate credentials.

Enter the number of seconds in the **Authorization refresh time** field. The Authorization Refresh Time allows you to manage how often the authorization data is verified with the authentication realm. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

4. In the **Inactivity timeout** field, enter the number of seconds to specify the amount of time a session can be inactive before it is logged out.
5. Configure cookie options:
 - a. Select the **Use persistent cookies** check box to use persistent browser cookies instead of session browser cookies.
 - b. Select the **Verify the IP address in the cookie** check box if you would like the cookies surrogate credentials to only be accepted for the IP address that the cookie was authenticated. Disabling this will allow cookies to be accepted from other IP addresses.
6. You can specify a virtual URL. For more information on the virtual URL, see ["About Origin-Style Redirection"](#) on page 910.
7. Select the **Challenge user after logout** check box if the realm requires the users to enter their credentials after they have logged out.
8. Click **Apply**.

Notes

If you use guest authentication/authorization:

- ❑ Local realms provide split authorization, and it is possible to be successfully authenticated but have authorization fail.
- ❑ If the Local realm `validate authorized user` command is disabled and the user does not exist in the authorization realm, authorization is considered a success and the user is assigned to the default group if there is one configured and it is of interest to policy.

Defining the Local User List

Defining the local user list involves the following steps:

- ❑ Create a list or customize the default list for your needs.
- ❑ Upload a user list or add users and groups through the CLI.
- ❑ Associate the list with the realm.

Creating a Local User List

The user list `local_user_database` is created on a new system or after an upgrade. It is empty on a new system. If a password file existed on the ProxySG before an upgrade, then the list contains all users and groups from the password file; the initial default user list is `local_user_database`. If a new user list is created, the

default can be changed to point to it instead by invoking the `security local-user-list default list list_name` command. You can create up to 50 new lists with 10,000 users each.

Lists can be uploaded or you can directly edit lists through the CLI. If you want to upload a list, it must be created as a text file using the `.htpasswd` format of the ProxySG.

Each user entry in the list consists of:

- ❑ username
- ❑ List of groups
- ❑ Hashed password
- ❑ Enabled/disabled boolean searches

A list that has been populated looks like this:

```
SGOS#(config) security local-user-list edit list_name
SGOS#(config local-user-list list_name) view
list20
Lockout parameters:
  Max failed attempts: 60
  Lockout duration:    3600
  Reset interval:      7200
Users:
admin1
  Hashed Password: $1$TvEzpzE$Z2A/OuJU3w5LnEONDHkmg.
  Enabled: true
  Groups:
    group1
admin2
  Hashed Password: $1$sKJvNB3r$xsInBU./2hhBz6xDAHpND.
  Enabled: true
  Groups:
    group1
    group2
admin3
  Hashed Password: $1$duuCUt30$keSdIkZVS4RyFz47G78X20
  Enabled: true
  Groups:
    group2
Groups:
  group1
  group2
```

To create a new empty local user list:

```
SGOS#(config) security local-user-list create list_name
```

Username

The username must be case-sensitively unique, and can be no more than 64 characters long. All characters are valid, except for a colon (:).

A new local user is enabled by default and has an empty password.

List of Groups

You cannot add a user to a group unless the group has previously been created in the list. The group name must be case-sensitively unique, and can be no more than 64 characters long. All characters are valid, except for colon (:).

The groups can be created in the list; however, their user permissions are defined through policies only.

Hashed Password

The hashed password must be a valid UNIX DES or MD5 password whose plain-text equivalent cannot be more than 64 characters long.

To populate the local user list using an off-box `.htpasswd` file, continue with the next section. To populate the local user list using the ProxySG CLI, go to ["Defining the Local User List"](#) on page 936.

Populating a List using the `.htpasswd` File

To add users to a text file in `.htpasswd` format, enter the following UNIX `htpasswd` command:

```
prompt> htpasswd [-c] .htpasswd username
```

The `-c` option creates a new `.htpasswd` file and should only be used for the very first `.htpasswd` command. You can overwrite any existing `.htpasswd` file by using the `-c` option.

After entering this command, you are prompted to enter a password for the user identified by `username`. The entered password is hashed and added to the user entry in the text file. If the `-m` option is specified, the password is hashed using MD5; otherwise, UNIX DES is used.

Important: Because the `-c` option overwrites the existing file, do not use the option if you are adding users to an existing `.htpasswd` file.

After you add the users to the `.htpasswd` file, you can manually edit the file to add user groups. When the `.htpasswd` file is complete, it should have the following format:

```
user:encrypted_password:group1,group2,...
user:encrypted_password:group1,group2,...
```

Note: You can also modify the users and groups once they are loaded on the ProxySG. To modify the list once it is on the appliance, see ["Populating a Local User List through the ProxySG"](#) on page 939.

Uploading the `.htpasswd` File

When the `.htpasswd` file is uploaded, the entries from it either replace all entries in the default local user list or append to the entries in the default local user list. One default local user list is specified on the ProxySG.

To set the default local user list use the command `security local-user-list default list list_name`. The list specified must exist.

To specify that the uploaded `.htpasswd` file replace all existing user entries in the default list, enter `security local-user-list default append-to-default disable` before uploading the `.htpasswd` file.

To specify that the `.htpasswd` file entries should be appended to the default list instead, enter `security local-user-list default append-to-default enable`.

To upload the `.htpasswd` file:

The `.htpasswd` file is loaded onto the ProxySG with a Perl script. The *SGOS 6.4 Release Notes* provide the current location for this file.

Unzip the file, which contains the `set_auth.pl` script.

Note: To use the `set_auth.pl` script, you must have Perl binaries on the system where the script is running.

To load the `.htpasswd` file:

```
prompt> set_auth.pl username password
path_to_.htpasswd_file_on_local_machine ip_address_of_the_ProxySG

where username and password are valid administrator credentials for the
ProxySG.
```

Populating a Local User List through the ProxySG

You can populate a local user list from scratch or modify a local user list that was populated by loading an `.htpasswd` file.

To create a new, empty local user list:

```
SGOS#(config) security local-user-list create list_name
```

To modify an existing local user list (can be empty or contain users):

- ❑ To enter configuration mode:

```
SGOS#(config) security local-user-list edit list_name
SGOS#(config local-user-list list_name)
```

- ❑ The following subcommands are available:

Note: To add users and groups to the list, enter the following commands, beginning with groups, since they must exist before you can add them to a user account.

```
SGOS#(config local-user-list list_name) group create group1
SGOS#(config local-user-list list_name) group create group2
SGOS#(config local-user-list list_name) group create group3
SGOS#(config local-user-list list_name) user create username
SGOS#(config local-user-list list_name) user edit username
SGOS#(config local-user-list list_name username) group add groupname1
```

```
SGOS#(config local-user-list list_name username) group add groupname2
SGOS#(config local-user-list list_name username) password password
-or-
SGOS#(config local-user-list list_name username) hashed-password
hashed-password
```

Note: If you enter a plain-text password, the ProxySG hashes the password. If you enter a hashed password, the appliance does not hash it again.

1. (Optional) The user account is enabled by default. To disable a user account:

```
SGOS#(config local-user-list list_name username) disable
ok
```

2. Repeat for each user you want added to the list.

To view the results of an individual user account:

Remain in the user account submode and enter the following command:

```
SGOS#(config local-user-list list_name username) view
admin1
  Hashed Password: $1$TvEzpZE$Z2A/OuJU3w5LnEONDHkmg.
  Enabled: true
  Failed Logins: 6
  Groups:
    group1
```

Note: If a user has no failed logins, the statistic does not display.

To view the users in the entire list:

Exit the user account submode and enter:

```
SGOS#(config local-user-list list_name username) exit
SGOS#(config local-user-list list_name) view
list20
Lockout parameters:
  Max failed attempts: 60
  Lockout duration:    3600
  Reset interval:     7200
Users:
admin1
  Hashed Password: $1$TvEzpZE$Z2A/OuJU3w5LnEONDHkmg.
  Enabled: true
  Groups:
    group1
admin2
  Hashed Password: $1$sKJvNB3r$xsInBU./2hhBz6xDAHpND.
  Enabled: true
  Groups:
    group1
    group2
admin3
  Hashed Password: $1$duuCUt30$keSdIkZVS4RyFz47G78X20
  Enabled: true
```

```
Groups:
  group2
Groups:
  group1
  group2
```

To view all the lists on the ProxySG:

```
SGOS#(config) show security local-user-list
Default List: local_user_database
Append users loaded from file to default list: false
local_user_database
Lockout parameters:
  Max failed attempts: 60
  Lockout duration:    3600
  Reset interval:     7200
Users:
  Groups:
test1
  Users:
  Groups:
```

To delete groups associated with a user:

```
SGOS#(config local-user-list list_name username) group remove
group_name
```

To delete users from a list:

```
SGOS#(config local-user-list list_name) user delete username
This will permanently delete the object. Proceed with deletion?
(y or n) y
ok
```

To delete all users from a list:

```
SGOS#(config local-user-list list_name) user clear
ok
```

The groups remain but have no users.

To delete all groups from a list:

```
SGOS#(config local-user-list list_name) group clear
ok
```

The users remain but do not belong to any groups.

Enhancing Security Settings for the Local User List

You can configure a local user database so that each user account is automatically disabled if too many failed login attempts occur for the account in too short a period, indicating a brute-force password attack on the ProxySG. The security settings are available through the CLI only.

Available security settings are:

- ❑ **Maximum failed attempts:** The maximum number of failed password attempts allowed for an account. When this threshold is reached, the account is disabled (locked). If this is zero, there is no limit. The default is 60 attempts.

- ❑ Lockout duration: The time after which a locked account is re-enabled. If this is zero, the account does not automatically re-enable, but instead remains locked until manually enabled. The default is 3600 seconds (one hour).
- ❑ Reset interval: The time after which a failed password count resets after the last failed password attempt. If this is zero, the failed password count resets only when the account is enabled or when its password is changed. The default is 7200 seconds (two hours).

These values are enabled by default on the system for all user account lists. You can change the defaults for each list that exists on the system.

To change the security settings for a specific user account list:

1. Enter the following commands from the (config) prompt:

```
SGOS#(config) security local-user-list edit list_name
SGOS#(config local-user-list list_name) lockout-duration seconds
SGOS#(config local-user-list list_name) max-failed-attempts attempts
SGOS#(config local-user-list list_name) reset-interval seconds
```

2. (Optional) View the settings:

```
SGOS#(config local-user-list list_name) view
listname
Lockout parameters:
  Max failed attempts: 45
  Lockout duration:    3600
  Reset interval:      0
```

3. (Optional) To disable any of these settings:

```
SGOS#(config local-user-list list_name) no [lockout-duration | max-
failed-attempts | reset-interval]
```

Creating the CPL

Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes. (The default policy in these examples is deny.)

Note: Refer to the *SGOS 6.4 Content Policy Language Reference* for details about CPL and how transactions trigger the evaluation of policy file layers.

- ❑ Every Local-authenticated user is allowed access the ProxySG.

```
<Proxy>
  authenticate(LocalRealm)
```

- ❑ Group membership is the determining factor in granting access to the ProxySG.

```
<Proxy>
  authenticate(LocalRealm)
<Proxy>
  group="group1" allow
```

- ❑ A subnet definition determines the members of a group, in this case, members of the Human Resources department.

```
<Proxy>
    authenticate(LocalRealm)
<Proxy>
    Define subnet HRSubnet
        192.168.0.0/16
        10.0.0.0/24
    End subnet HRSubnet
    [Rule] client_address=HRSubnet
        url.domain=monster.com
        url.domain=hotjobs.com
        deny
.
.
.
    [Rule]
        deny
```


Chapter 48: CA eTrust SiteMinder Authentication

The ProxySG can be configured to consult a SiteMinder policy server for authentication and session management decisions. This requires that a SiteMinder realm be configured on the ProxySG and policy written to use that realm for authentication.

Access to the SiteMinder policy server is done through the Blue Coat Authentication and Authorization Agent (BCAAA).

SiteMinder authentication cannot be used to authenticate administrative users to the ProxySG appliance management console.

Note: Refer to the *BCAAA Service Requirements* document for up-to-date information on BCAA compatibility and installation. The *BCAAA Service Requirements* document is posted at MySymantec.

Topics in this Section

This section includes information about the following topics:

- ❑ ["About SiteMinder Interaction with Blue Coat"](#) on page 945
- ❑ ["Participating in a Single Sign-On \(SSO\) Scheme"](#) on page 948
- ❑ ["Creating a SiteMinder Realm"](#) on page 949
- ❑ ["Configuring SiteMinder Servers"](#) on page 951
- ❑ ["Defining SiteMinder Server General Properties"](#) on page 952
- ❑ ["Creating the CPL"](#) on page 957
- ❑ ["SiteMinder Authorization Example"](#) on page 958

About SiteMinder Interaction with Blue Coat

Within the SiteMinder system, BCAA acts as a custom Web agent. It communicates with the SiteMinder policy server to authenticate the user and to obtain a SiteMinder session token, response attribute information, and group membership information.

Custom header and cookie response attributes associated with **OnAuthAccept** and **OnAccessAccept** attributes are obtained from the policy server and forwarded to the ProxySG. They can (as an option) be included in requests forwarded by the appliance.

Within the ProxySG system, BCAA acts as its agent to communicate with the SiteMinder server. The ProxySG provides the user information to be validated to BCAA, and receives the session token and other information from BCAA.

Each ProxySG SiteMinder realm used causes the creation of a BCAA process on the Windows or Solaris host computer running BCAA. A single host computer can support multiple ProxySG realms (from the same or different ProxySG appliances); the number depends on the capacity of the BCAA host computer and the amount of activity in the realms.

Note: Each (active) SiteMinder realm on the ProxySG must reference a different agent on the Policy Server.

Configuration of the ProxySG realm must be coordinated with configuration of the SiteMinder policy server. Each must be configured to be aware of the other. In addition, certain SiteMinder responses must be configured so that BCAA gets the information the ProxySG needs.

Configuring the SiteMinder Policy Server

Note: Blue Coat assumes you are familiar with configuration of SiteMinder policy servers and Web agents.

Since BCAA is a Web agent in the SiteMinder system, it must be configured on the SiteMinder policy server. Configuration of BCAA on the host computer is not required; the agent obtains its configuration information from the ProxySG.

A suitable Web agent must be created and configured on the SiteMinder server. This must be configured to support 5.x agents, and a shared secret must be chosen and entered on the server (it must also be entered in the ProxySG SiteMinder realm configuration).

SiteMinder protects resources identified by URLs. A ProxySG realm is associated with a single protected resource. This could be an already existing resource on a SiteMinder server, (typical for a reverse proxy arrangement) or it could be a resource created specifically to protect access to ProxySG services (typical for a forward proxy).

Note: The request URL is not sent to the SiteMinder policy server as the requested resource; the requested resource is the entire ProxySG realm. Access control of individual URLs is done on the ProxySG using CPL or VPM.

The SiteMinder realm that controls the protected resource must be configured with a compatible authentication scheme. The supported schemes are Basic (in plain text and over SSL), Forms (in plain text and over SSL), and X.509 certificates. Configure the SiteMinder realm with one of these authentication schemes.

Note: Only the following X.509 Certificates are supported: X.509 Client Cert Template, X.509 Client Cert and Basic Template, and X.509 Client Cert and Form Template.

The ProxySG requires information about the authenticated user to be returned as a SiteMinder response. The responses should be sent by an `OnAuthAccept` rule used in the policy that controls the protected resource.

The responses must include the following:

- ❑ A Web-Agent-HTTP-Header-variable named `BCSI_USERNAME`. It must be a user attribute; the value of the response must be the simple username of the authenticated user. For example, with an LDAP directory this might be the value of the `cn` attribute or the `uid` attribute.
- ❑ A Web-Agent-HTTP-Header-variable named `BCSI_GROUPS`. It must be a user attribute and the value of the response must be `SM_USERGROUPS`.

If the policy server returns an LDAP FQDN as part of the authentication response, the ProxySG uses that LDAP FQDN as the FQDN of the user.

Once the SiteMinder agent object, configuration, realm, rules, responses and policy have been defined, the ProxySG can be configured.

Additional SiteMinder Configuration Notes

Note: Additional configuration might be needed on the SiteMinder server depending on specific features being used.

- ❑ If using single-sign on (SSO) with off-box redirection (such as to a forms login page), the forms page must be processed by a 5.x or later Web Agent, and that agent must be configured with `fccompatmode=no`. This keeps that agent from doing SSO with 5.x agents.
- ❑ For SSO to work with other Web agents, the other agents must have the `AcceptTPCookie=YES` as part of their configuration. This is described in the SiteMinder documentation.
- ❑ Blue Coat does not extract the issuerDN from X.509 certificates in the same way as the SiteMinder agent. Thus, a separate certificate mapping might be needed for the SGOS agent and the SiteMinder agents.

For example, the following was added to the SiteMinder policy server certificate mappings:

```
CN=Waterloo Authentication and Security Team,OU=Waterloo R&D, O=Blue  
Coat\, Inc.,L=Waterloo,ST=ON,C=CA
```

- ❑ In order to use off-box redirection (such as an SSO realm), all agents involved must have the setting `EncryptAgentName=no` in their configurations.
- ❑ The ProxySG's credential cache only caches the user's authentication information for the smaller of the time-to-live (TTL) configured on the ProxySG and the session TTL configured on the SiteMinder policy server.

Configuring the ProxySG Realm

The ProxySG realm must be configured so that it can:

- ❑ Find the BCAAA service that acts on its behalf (hostname or IP address, port, SSL options, and the like).
- ❑ Provide BCAAA with the information necessary to allow it to identify itself as a Web agent (agent name, shared secret).
- ❑ Provide BCAAA with the information that allows it to find the SiteMinder policy server (IP address, ports, connection information.)
- ❑ Provide BCAAA with the information that it needs to do authentication and collect authorization information (protected resource name), and general options (server fail-over and off-box redirection)

For more information on configuring the ProxySG SiteMinder realm, see ["Creating a SiteMinder Realm"](#) on page 949.

Note: All ProxySG and agent configuration occurs on the appliance. The ProxySG sends the necessary information to BCAAA when it establishes communication.

Participating in a Single Sign-On (SSO) Scheme

The ProxySG can participate in SSO with other systems that use the same SiteMinder policy server. Users must supply their authentication credentials only once to any of the systems participating. Participating in SSO is not a requirement, the ProxySG can use the SiteMinder realm as an ordinary realm.

When using SSO with SiteMinder, the SSO token is carried in a cookie (`SMSESSION`). This cookie is set in the browser by the first system that authenticates the user; other systems obtain authentication information from the cookie and so do not have to challenge the user for credentials. The ProxySG sets the `SMSESSION` cookie if it is the first system to authenticate a user, and authenticates the user based on the cookie if the cookie is present.

Since the SSO information is carried in a cookie, all the servers participating must be in the same cookie domain, including the ProxySG. This imposes restrictions on the `authenticate.mode()` used on the ProxySG.

- ❑ A reverse proxy can use any `origin` mode.
- ❑ A forward proxy must use one of the `origin-redirect` modes (such as `origin-cookie-redirect`). When using `origin-*-redirect` modes, the virtual URL hostname must be in the same cookie domain as the other systems. It cannot be an IP address; the default `www.cfauth.com` does not work either.

When using `origin-*-redirect`, the SSO cookie is automatically set in an appropriate response after the ProxySG authenticates the user. When using `origin` mode (in a reverse proxy), setting this cookie must be explicitly specified by the administrator. The policy substitution variable `$(x-agent-sso-cookie)` expands to the appropriate value of the `set-cookie:` header.

Avoiding ProxySG Challenges

In some SiteMinder deployments all credential challenges are issued by a central authentication service (typically a Web server that challenges through a form). Protected services do not challenge and process request credentials; instead, they work entirely with the SSO token. If the request does not include an SSO token, or the SSO token is not acceptable, the request is redirected to the central service, where authentication occurs. After authentication completes, the request redirects to the original resource with a response that sets the SSO token.


If the SiteMinder policy server is configured to use a forms-based authentication scheme, the above happens automatically. However, in this case, the ProxySG realm can be configured to redirect to an off-box authentication service always. The URL of the service is configured in the scheme definition on the SiteMinder policy server. The ProxySG realm is then configured with `always-redirect-offbox` enabled.

The ProxySG must not attempt to authenticate a request for the off-box authentication URL. If necessary, `authenticate(no)` can be used in policy to prevent this.

Creating a SiteMinder Realm

To create a SiteMinder realm:

1. Select the **Configuration > Authentication > CA eTrust SiteMinder > SiteMinder Realms** tab.
2. Click **New**. The Add SiteMinder Realm dialog displays.



3. In the **Realm name** field, enter a realm name. The name can be 32 characters long and composed of alphanumeric characters and underscores. The name *must* start with a letter. The name should be meaningful to you, but it does not have to be the name of the SiteMinder policy server.
4. Click **OK**.
5. Click **Apply**.

Configuring SiteMinder Agents

You must configure the SiteMinder realm so that it can find the Blue Coat Authentication and Authorization Agent (BCAAA).

1. Select the **Configuration > Authentication > CA eTrust SiteMinder > Agents** tab.

The screenshot shows the 'Agents' configuration window for SiteMinder. It includes fields for Realm name, Primary agent (Host, Port, Agent name, Change Secret), Alternate agent (Host, Port, Agent name, Change Secret), SSL Options (Enable SSL, SSL device profile), and Timeout request after (60 seconds). A 'Case sensitive' checkbox is also present.


2. Select the realm name to edit from the drop-down list.
3. Configure the primary agent:
 - a. In the **Primary** agent section, enter the hostname or IP address where the agent resides.
 - b. Change the port from the default of **16101** if your SiteMinder port is different.
 - c. Enter the agent name in the **Agent name** field. The agent name is the name as configured on the SiteMinder policy server.
 - d. You must create a secret for the Agent that matches the secret created on the SiteMinder policy server. Click **Change Secret**. SiteMinder secrets can be up to 64 characters long and are always case sensitive.
4. (Optional) Enter an alternate agent host and agent name in the **Alternate agent** section.
5. Configure SSL options:
 - a. (Optional) Click **Enable SSL** to enable SSL between the ProxySG and the BCAAA service.
 - b. (Optional) Select the SSL device profile that this realm uses to make an SSL connection to a remote system. You can choose any device profile that displays in the drop-down list. For information on using device profiles, see "[Appliance Certificates and SSL Device Profiles](#)" on page 1290.
6. In the **Timeout Request** field, enter the number of seconds the ProxySG allows for each request attempt before timing out. (The default request timeout is **60** seconds.)

7. If you want group comparisons for SiteMinder groups to be case sensitive, select **Case sensitive**.
8. Click **Apply**.


Configuring SiteMinder Servers

Once you create a SiteMinder realm, use the SiteMinder Servers page to create and edit the list of SiteMinder policy servers consulted by the realm.

1. Select the **Configuration > Authentication > CA eTrust SiteMinder > SiteMinder Servers** tab.
2. From the **Realm name** drop-down list, select the SiteMinder realm for which you want to add servers or change server properties.
3. To create a new SiteMinder policy server, click **New**. The Add List Item dialog displays.



4. Enter the name of the server in the dialog. This name is used only to identify the server in the ProxySG's configuration; it usually is the real hostname of the SiteMinder policy server.
5. Click **OK**.
6. To edit an existing SiteMinder policy server, highlight the server and click **Edit**. The Edit SiteMinder Server.



7. Configure the server options:
 - a. Enter the IP address of the SiteMinder policy server in the **IP address** field.
 - b. Enter the correct port numbers for the **Authentication**, **Authorization**, and **Accounting** ports, which are the same ports configured on their SiteMinder policy server. The valid port range is 1-65535.
 - c. The **Maximum Connections** to the server is 32768; the default is **256**.
 - d. The **Connection Increment** specifies how many connections to open at a time if more are needed and the maximum has not been exceeded. The default is **1**.
 - e. The **Timeout** value has a default of **60** seconds, which can be changed.
 - f. Click **OK**.
8. Click **Apply**.

Defining SiteMinder Server General Properties

The **SiteMinder Server General** tab allows you to specify the protected resource name, the server mode, and whether requests should always be redirected off box.

To configure general settings:

1. Select the **Configuration > Authentication > CA eTrust SiteMinder > SiteMinder Server General** tab.

2. Configure the following options:
 - a. From the **Realm name** drop-down list, select the SiteMinder realm for which you want to change properties.
 - b. Enter the **Protected resource name**. The protected resource name is the same as the resource name on the SiteMinder policy server that has rules and policy defined for it. When entering a protected resource name, precede it with a forward slash (/). For example, if the protected resource name is `bcsi`, you would enter `/bcsi`.
 - c. In the **Server mode** drop-down list, select either **failover** or **round-robin**. Failover mode falls back to one of the other servers if the primary one is down. Round-robin modes specifies that all of the servers should be used together in a round-robin approach. Failover is the default.

Note: The server mode describes the way the agent (the BCAAA service) interacts with the SiteMinder policy server, not the way that ProxySG interacts with BCAAA.

3. To force authentication challenges to always be redirected to an off-box URL, select **Always redirect off-box**.

Note: All SiteMinder Web agents involved must have the setting `EncryptAgentName=no` in their configurations to go off-box for any reason.

If using SiteMinder forms for authentication, the ProxySG always redirects the browser to the forms URL for authentication. You can force this behavior for other SiteMinder schemes by configuring the **always redirect off-box** property on the realm.

4. If your Web applications need information from the SiteMinder policy server responses, you can select **Add Header Responses**. Responses from the policy server obtained during authentication are added to each request forwarded by the ProxySG. Header responses replace any existing header of the same name; if no such header exists, the header is added. Cookie responses replace a cookie header with the same cookie name; if no such cookie header exists, one is added.

5. To enable validation of the client IP address, select **Validate client IP address**. If the client IP address in the SSO cookie can be valid yet different from the current request client IP address, due to downstream proxies or other devices, clear the **Validate client IP address** option for the realm. Also modify the SiteMinder agents participating in SSO with the ProxySG; set the **TransientIPCheck** variable to **yes** to enable IP address validation and **no** to disable it.
6. Click **Apply**.

Configuring Authorization Settings for SiteMinder

The **Authorization** tab allows you to authorize users through another realm and specify search criteria for the user ID.

To specify authorization settings for SiteMinder:

1. Select the **Configuration > Authentication > CA eTrust SiteMinder > Authorization** tab.

2. From the **Realm** name drop-down list, select a SiteMinder realm.
3. From the **Authorization realm name** drop-down list, select the LDAP, Local, or XML realm you want to use to authorize users. If **Self** is selected, the **Authorization username** must be **Use FQDN**.
4. Configure authorization options. You cannot always construct the user's authorization username from the substitutions available. If not, you can search on a LDAP server for a user with an attribute matching the substitution and then use the FQDN for the matched user as the authorization username. Authorization then occurs on that authorization username:
 - a. In the **Authorization username** field, enter the substitution to use to identify the user. The default authorization username is `$(cs-username)`. You can use any policy substitutions. -or-

- b. Select **Use FQDN** or to determine through search criteria, which uses the FQDN or full username determined while identifying the user during the authentication process. -or-
- c. Select **Determine by search**, which enables the fields below. Specify the following to focus the search:
 - **LDAP search realm name:** An LDAP realm to search. In most cases, this is the same as the LDAP realm used for authorization.
 - **Search filter:** Used during the LDAP search. This search filter can contain policy substitutions including the \$(cs-username) substitution.
 - **User attribute:** An attribute on the entry returned in the LDAP search results that has the value to use as the authorization username. In most cases this is the FQDN of the user entry.
5. (Optional) Click **Set Users to Ignore** to add a list of users excluded from searches.
6. Click **Apply**.

Configuring General Settings for SiteMinder

The SiteMinder General tab allows you to specify a display name, the refresh times, a inactivity timeout value, cookies, and a virtual URL.

To configure general settings for SiteMinder:

1. Select the **Configuration > Authentication > CA eTrust SiteMinder > SiteMinder General** tab.

The screenshot shows the 'SiteMinder General' configuration tab. On the left, there are numbered callouts (2 through 9) with arrows pointing to specific settings. The settings are as follows:

- 2** points to the **Realm name:** dropdown menu, which is set to 'SM_test'.
- 3** points to the **Display name:** text field, which contains 'SM_test'.
- 4** points to the **Refresh Times:** section, which includes a checked checkbox for 'Use the same refresh time for all', and two text fields: 'Credential refresh time: 900 seconds' and 'Surrogate refresh time: 900 seconds'.
- 5** points to the **Inactivity timeout:** text field, which contains '900 seconds'.
- 6** points to the **Rejected credentials time:** text field, which contains '1 seconds'.
- 7** points to the **Cookies** section, which includes two checkboxes: 'Use persistent cookies' (unchecked) and 'Verify the IP address in the cookie' (checked).
- 8** points to the **Virtual URL:** text field, which contains 'www.cfauth.com/'.
- 9** points to the **Challenge user after logout** checkbox, which is checked.

2. From the **Realm name** drop-down list, select the SiteMinder realm for which you want to change properties.

3. If needed, change the SiteMinder realm display name. The default value for the display name is the realm name. The display name cannot be greater than 128 characters and it cannot be empty.
4. Configure refresh options:
 - a. Select the **Use the same refresh time for all** check box if you would like to use the same refresh time for all.
 - b. Enter the number of seconds in the **Credential refresh time** field. The Credential Refresh Time is the amount of time Basic credentials (username and password) are kept on the ProxySG. This feature allows the ProxySG to reduce the load on the authentication server and enables credential spoofing. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here. Before the refresh time expires, the ProxySG authenticates the user supplied credentials against the cached credentials. If the credentials received do not match the cached credentials, they are forwarded to the authentication server in case the user password changed. After the refresh time expires, the credentials are forwarded to the authentication server for verification.
 - c. Enter the number of seconds in the **Surrogate refresh time** field. The Surrogate Refresh Time allows you to set a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the ProxySG authenticates the transaction. After the refresh time expires, the ProxySG verifies the user's credentials. Depending upon the authentication mode and the user-agent, this may result in challenging the end user for credentials.

The main goal of this feature is to verify that the user-agent still has the appropriate credentials.
5. Enter the number of seconds in the **Inactivity timeout** field to specify the amount of time a session can be inactive before being logged out.
6. If you use Basic credentials and want to cache failed authentication attempts (to reduce the load on the authentication service), enter the number of seconds in the **Rejected Credentials time** field. This setting, enabled by default and set to one second, allows failed authentication attempts to be automatically rejected for up to 10 seconds. Any Basic credentials that match a failed result before its cache time expires are rejected without consulting the back-end authentication

service. The original failed authentication result is returned for the new request. All failed authentication attempts can be cached: Bad password, expired account, disabled account, old password, server down. To disable caching for failed authentication attempts, set the **Rejected Credentials time** field to 0.

7. Configure cookie options:
 - a. Select the **Use persistent cookies** check box to use persistent browser cookies instead of session browser cookies.
 - b. Select the **Verify the IP address in the cookie** check box if you would like the cookies surrogate credentials to only be accepted for the IP address that the cookie was authenticated. Disabling this allows cookies to be accepted from other IP addresses.
8. Specify the virtual URL to redirect the user to when they need to be challenged by the ProxySG. If the appliance is participating in SSO, the virtual hostname must be in the same cookie domain as the other servers participating in the SSO. It cannot be an IP address or the default, `www.cfauth.com`.
9. Select the **Challenge user after logout** check box if the realm requires the users to enter their credentials after they have logged out.
10. Click **Apply**.

Creating the CPL

You can create CPL policies now that you have completed SiteMinder realm configuration. Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes.

The examples below assume the default policy condition is *allow*. On new SGOS systems, the default policy condition is *deny*.

Note: Refer to *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file `<Proxy>` and other layers.

- ❑ Every SiteMinder-authenticated user is allowed access the ProxySG.

```
<Proxy>
  authenticate (SiteMinderRealm)
```

- ❑ Group membership is the determining factor in granting access to the ProxySG.

```
<Proxy>
  authenticate (SiteMinderRealm)
<Proxy>
  group="cn=proxyusers, ou=groups, o=myco"
  deny
```

SiteMinder Authorization Example

Situation

Credential challenges are issued by a central authentication service (this means the SiteMinder realm must be enabled to always redirect authentication requests offbox), and an LDAP search can be used to find the FQDN.

Configuration

1. Download and install the BCAA service.
2. Set up the SiteMinder server; be sure to configure the SMSession cookie and the BCSI_USERNAME variable on the SiteMinder server.
3. Configure an LDAP, XML, or Local realm that can be used to authorize users.
4. Create and define a ProxySG SiteMinder realm. Specifically:
 - Use the **Agents** tab to configure the BCAA service and the SiteMinder service to work with the SiteMinder server.
 - Use the **SiteMinder Server** tab to associate the realm with the SiteMinder server.
 - Use the **SiteMinder Server General** tab to always redirect requests off box.
 - Use the **Authorization** tab to set up search criteria for user IDs.

Behavior

- ❑ ProxySG receives a request for a user.
 - If this request does not contain an SMSession cookie (user unauthenticated), the ProxySG redirects the request to the central authentication service. The URL of the service is configured in the scheme definition on the SiteMinder policy server. When the request returns from the central authentication service, the SMSession cookie is extracted and sent to the BCAA service for validation.
 - If the request does contain an SMSession cookie, the ProxySG passes the SMSession cookie through the BCAA service for validation and authentication.
- ❑ The SiteMinder policy server authenticates the user and sends the LDAP attribute of the user (UID) in the BCSI_USERNAME variable to the BCAA service, which then passes it on the ProxySG.
- ❑ The ProxySG uses the UID attribute to do an LDAP search, identifying the user FQDN.
- ❑ The ProxySG uses the FQDN to construct an LDAP query to the authorization LDAP realm server to compare and validate group membership.

You can use the result to check group-based policy.

Chapter 49: Certificate Realm Authentication

If you have a Public Key Infrastructure (PKI) in place, you can configure the ProxySG to authenticate users based on their X.509 certificates by creating a certificate realm. Additionally, if the users are members of an LDAP, XML, or Local group, you can configure the certificate realm to forward the user credentials to the LDAP, XML, or Local realm for authorization.

X.509 Certificate authentication realms can be used to authenticate administrative users (read only and read/write) to the management console. To ensure that credentials are not sent in clear text, configure the Certificate realm to use TLS to secure the communication with the authorization server.

The following topics describe how to set up and configure a certificate realm:

- ❑ ["How a Certificate Realm Works"](#) on page 959
- ❑ ["Configuring Certificate Realms"](#) on page 960
- ❑ ["Specifying an Authorization Realm"](#) on page 964
- ❑ ["Revoking User Certificates"](#) on page 966
- ❑ ["Creating a Certificate Authorization Policy"](#) on page 966
- ❑ ["Tips"](#) on page 967
- ❑ ["Certificate Realm Example"](#) on page 968

How a Certificate Realm Works

After an SSL session has been established, the user is prompted to select the certificate to send to the ProxySG. If the certificate was signed by a Certificate Authority (CA) that the ProxySG trusts, the user is considered authenticated. The ProxySG then extracts the username for that user from the certificate.

At this point the user is authenticated. If an authorization realm has been specified, such as LDAP, XML or Local, the certificate realm then passes the username to the specified authorization realm, which figures out which groups the user belongs to.

Note: If you authenticate with a certificate realm, you cannot also challenge for a password.

Certificate realms do not require an authorization realm. If no authorization realm is configured, the user cannot be a member of any group. You do not need to specify an authorization realm if:

- ❑ The policy does not make any decisions based on groups
- ❑ The policy works as desired when all certificate realm-authenticated users are not in any group

Configuring Certificate Realms


To configure a certificate realm, you must:

- ❑ Configure SSL between the client and ProxySG. See ["Using SSL with Authentication and Authorization Services"](#) on page 919 for more information.
- ❑ Enable **verify-client** on the HTTPS reverse proxy service to be used. See ["Creating an HTTPS Reverse Proxy Service"](#) on page 321 for more information.
- ❑ Verify that the certificate authority that signed the client's certificates is in the ProxySG *trusted* list. See ["Importing CA Certificates"](#) on page 1141.
- ❑ Create the certificate realm as described in ["Creating a Certificate Realm"](#) on page 960.
- ❑ Specify the fields to extract from the client certificate as described in ["Configuring Certificate Realm Properties"](#) on page 961.
- ❑ Customize the certificate realm properties as described in ["Defining General Certificate Realm Properties"](#) on page 963.
- ❑ (optional) If you want to authorize users who are part of an LDAP, XML, or Local group, configure authorization as described in ["Specifying an Authorization Realm"](#) on page 964.

Creating a Certificate Realm

To create a certificate realm:

1. Select the **Configuration > Authentication > Certificate > Certificate Realms** tab.
2. Click **New**. The Add Certificate Realm dialog displays.



3. In the **Realm name** field, enter a realm name. The name can be 32 characters long and composed of alphanumeric characters and underscores. The name *must* start with a letter.
4. Click **OK**.
5. Click **Apply**.

Configuring Certificate Realm Properties

The **Certificate Main** tab allows an administrator to define substitutions used to extract user data from a user certificate. The username and full username data can come from almost any field of a certificate; however, the username can only reference data within the field, not the field as a whole. Supported fields include: **serialNumber**, **issuer**, **subject**, **issuerAltName**, and **subjectAltName** fields.

To define certificate authentication properties:

1. Select the **Configuration > Authentication > Certificate > Certificate Main** tab.

2. From the **Realm name** drop-down list, select the Certificate realm for which you want to change realm properties.
3. In the **username** field, enter the substitution that specifies the common name in the subject of the certificate. **\$(CN.1)** is the default. Be aware that multiple attributes can be entered into the field to build complex substitutions.
4. (Optional) In the **Full Username** field, enter the substitutions used to construct the user's full username. For example, the user principal name (UPN) or LDAP distinguished name (DN). The field is empty by default.

The substitutions used to construct the username use the following parser format:

Parser Format

```
$( [attributename=] [field] [ .generalName[ .generalNameindex] ]
[ .attribute[ .attribute index] ] )
```

To see how the parser works, examine the client certificate example and the resulting substitutions in the table.

Client Certificate Example

```
subject: CN=John,OU=Auth,OU=Waterloo,O=BlueCoat
subjectAltName:
  -otherName: john.doe@bluecoat.com
  -otherName: john.doe@department.bluecoat.com
  -DN: CN=Doe,
      john,CN=Users,DC=internal,DC=cacheFlow,DC=com
```

	Parser Function	Format	Example
1.	Multiple instances of an attribute/general name results in an attribute/general name list.	\$(<i>attribute value</i>) and \$(<i>subjectAltName.general name value</i>) and \$(<i>issuerAltName.<general name value></i>)	\$(OU) = AuthWaterloo
2.	An individual instance of a multiple valued field is selected using its index (1-based). Works for attribute and general name fields.	\$(<i>field name.index#</i>)	\$(OU.2) = Waterloo
3.	The subjectAltName and issuerAltName fields support general name types that can be specified in the substitution. If multiple values of the same general name are found, all values will be substituted in a list. Supported general names types are: - otherName - email - DNS - dirName - URI - IP - RID	\$(<i>subjectAltName.<general name value></i>) or \$(<i>issuerAltName.<general name value></i>)	\$(subjectAltName.othername) = john.doe@bluecoat.com john.doe@department.bluecoat.com
4.	A modifier that enables LDAP style expansion of attributes.	<i>attribute name=</i>	\$(OU=subject.OU) = OU=Auth,OU=Waterloo

	Parser Function	Format	Example
5.	Text that is not part of a substitution is directly placed into the username.	$\$(any\ value),text$	$\$(OU),o=example$ becomes AuthWaterloo,o=example

Note: Starting in SGOS version 5.4, the username is no longer appended to the container attribute list. If you upgrade from a previous version, the existing substitutions are converted to the new parser, but may require a manual update.

5. Add or delete OIDs to enforce Extended Key Usage fields in a certificate. The list is empty by default. For example, to enforce a Microsoft Smart Card Logon OID, add a valid OID such as `1.3.6.1.4.1.311.20.2.2`.
6. Click **Apply** to complete the changes.

Defining General Certificate Realm Properties

The **Certificate General** tab allows you to specify the display name, the refresh times, an inactivity timeout value, cookies, and a virtual URL.

To configure certificate realm general settings:

1. Select the **Configuration > Authentication > Certificate > Certificate General** tab.

The screenshot displays the 'Certificate General' configuration tab. It includes the following fields and options:

- 2** Realm name: CertTest (dropdown menu)
- 3** Display name: CertTest (text field)
- 4** Refresh Times:
 - ☒ Use the same refresh time for all
 - Surrogate refresh time: 900 seconds
 - Authorization refresh time: 900 seconds
- 5** Inactivity timeout: 900 seconds
- 6** Cookies:
 - ☐ Use persistent cookies
 - ☒ Verify the IP address in the cookie
- 7** Virtual URL: www.cfauth.com/

2. From the **Realm name** drop-down list, select the Certificate realm to modify.
3. If necessary, change the realm's display name.
4. Configure refresh options:
 - a. Select **Use the same refresh time for all** to use the same refresh time for all.

- b. Enter the number of seconds in the **Surrogate refresh time** field. The Surrogate Refresh Time allows you to set a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the ProxySG authenticates the transaction. After the refresh time expires, the ProxySG will verify the user's certificate.
 - c. Enter the number of seconds in the **Authorization refresh time** field. The Authorization Refresh Time allows you to manage how often the authorization data is verified with the authentication realm. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.
5. Enter the number of seconds in the **Inactivity timeout** field to specify the amount of time a session can be inactive before being logged out.
6. Configure cookie options:
 - a. Select the **Use persistent cookies** check box to use persistent browser cookies instead of session browser cookies.
 - b. Select the **Verify the IP address in the cookie** check box if you would like the cookies surrogate credentials to only be accepted for the IP address that the cookie was authenticated. Disabling this allows cookies to be accepted from other IP addresses.
7. You can specify a virtual URL. For more information on the virtual URL, see ["About Origin-Style Redirection"](#) on page 910.
8. Click **Apply**.

Specifying an Authorization Realm

The Authorization tab allows you to set the authorization realm and to determine the authorization username.

To set certificate realm authorization properties:

1. Select the **Configuration > Authentication > Certificate > Authorization** tab.

2. Select the certificate realm for which you want to configure authorization from the **Realm name** drop-down list.
3. Select the realm that you will use for authorization from the **Authorization realm name** drop-down list. You can use an LDAP, Local, or XML realm to authorize the users in a certificate realm.
4. Configure authorization options. You cannot always construct the user's authorization username from the substitutions available. If not, you can search on a LDAP server for a user with an attribute matching the substitution and then use the FQDN for the matched user as the authorization username. Authorization would then be done on that authorization username.:
 - a. In the **Authorization username** field, enter the substitution to use to identify the user. The default authorization username is `$(cs-username)`. You can use any policy substitutions. -or-
 - b. Select **Use FQDN** or to determine through search criteria, which uses the FQDN or full username determined while identifying the user during the authentication process. -or-
 - c. Select **Determine by search**, which enables the fields below. Specify the following to focus the search:
 - **LDAP search realm name:** An LDAP realm to search. In most cases, this is the same as the LDAP realm used for authorization.
 - **Search filter:** Used during the LDAP search. This search filter can contain policy substitutions, including the `$(cs-username)` substitution.
 - **User attribute:** An attribute on the entry returned in the LDAP search results that has the value to use as the authorization username. In most cases this is the FQDN of the user entry.

5. (Optional) Click **Set Users to Ignore** to add a list of users excluded from searches.
6. Click **Apply**.

Revoking User Certificates

Using policy, you can revoke certain certificates by writing policy that denies access to users who have authenticated with a certificate you want to revoke. You must maintain this list on the ProxySG; it is not updated automatically.

Note: This method of revoking user certificates is meant for those with a small number of certificates to manage. For information on using automatically updated lists, see ["Using Certificate Revocation Lists"](#) on page 1134.

A certificate is identified by its issuer (the Certificate Signing Authority that signed it) and its serial number, which is unique to that CA.

Using that information, you can use the following strings to create a policy to revoke user certificates:

- ❑ `user.x509.serialNumber`—This is a string representation of the certificate's serial number in HEX. The string is always an even number of characters long, so if the number needs an odd number of characters to represent in hex, there is a leading zero. Comparisons are case insensitive.
- ❑ `user.x509.issuer`—This is an RFC2253 LDAP DN. Comparisons are case sensitive.
- ❑ (optional) `user.x509.subject`: This is an RFC2253 LDAP DN. Comparisons are case sensitive.

Example

If you have only one Certificate Signing Authority signing user certificates, you do not need to test the issuer. In the `<Proxy>` layer of the Local Policy file:

```
<proxy>
  deny user.x509.serialnumber=11
  deny user.x509.serialNumber=0F
```

If you have multiple Certificate Signing Authorities, test both the issuer and the serial number. In the `<Proxy>` layer of the Local Policy file:

```
<proxy>
  deny
  user.x509.issuer="Email=name,CN=name,OU=name,O=company,L=city,ST=state
  or province,C=country" user.x509.serialnumber=11\
  deny user.x509.issuer="CN=name,OU=name,O=company, L=city,ST=state or
  province,C=country" \
  deny user.x509.serialnumber=2CB06E9F000000000000B
```

Creating a Certificate Authorization Policy

When you complete Certificate realm configuration, you can create CPL policies. Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file `<Proxy>` and other layers.

Be aware that the default policy condition for these examples is *allow*. On new SGOS systems, the default policy condition is *deny*.

- ❑ Every certificate realm authenticated user is allowed access the ProxySG.

```
<Proxy>
    authenticate(CertificateRealm)
```

- ❑ A subnet definition determines the members of a group, in this case, members of the Human Resources department. (They are allowed access to the two URLs listed. Everyone else is denied permission.)

```
<Proxy>
    authenticate(CertificateRealm)
<Proxy>
    Define subnet HRSubnet
        192.168.0.0/16
        10.0.0.0/24
    End subnet HRSubnet
    [Rule] client_address=HRSubnet
        url.domain=monster.com
        url.domain=hotjobs.com
        deny
    .
    .
    .
    [Rule]
        deny
```

Tips

If you use a certificate realm and see an error message similar to the following

```
Realm configuration error for realm "cert": connection is not SSL.
```

This means that certificate authentication was requested for a transaction, but the transaction was not done on an SSL connection, so no certificate was available.

This can happen in three ways:

- ❑ The authenticate mode is either `origin-IP-redirect/origin-cookie-redirect` or `origin-IP/origin-cookie`, but the virtual URL does not have an `https:` scheme. This is likely if authentication through a certificate realm is selected with no other configuration, because the default configuration does not use SSL for the virtual URL.

- ❑ In a server accelerator deployment, the authenticate mode is origin and the transaction is on a non-SSL port.
- ❑ The authenticate mode is `origin-IP-redirect/origin-cookie-redirect`, the user has authenticated, the credential cache entry has expired, and the next operation is a POST or PUT from a browser that does not handle 307 redirects (that is, from a browser other than Internet Explorer). The workaround is to visit another URL to refresh the credential cache entry and then try the POST again.

Certificate Realm Example

Situation

Reverse proxy with user authentication and authorization from the ProxySG in combination with an LDAP server and an end-user PKI certificate. The subject of the certificate includes the e-mail address of the user.

Configuration

1. Configure an HTTPS reverse proxy as explained in ["Creating an HTTPS Reverse Proxy Service"](#) on page 321. Be sure to enable the **Verify Client** option.
2. Configure SSL between the client and ProxySG (for more information, see ["Using SSL with Authentication and Authorization Services"](#) on page 919).
3. Verify that the certificate authority that signed the client's certificates is in the ProxySG *trusted* list.
4. Make sure that ProxySG CRL is correct (for more information, see ["Using Certificate Revocation Lists"](#) on page 1134.)
5. Create a Certificate Authority Certificate List (CCL) and add the CA that created the certificate to the CCL. (For more information, see ["Managing CA Certificate Lists"](#) on page 1143.)
6. Configure the certificate realm:
 - Use the **Configuration > Authentication > Certificate > Realms** tab to name the realm.
 - Use the **Configuration > Authentication > Certificate > Main** tab to define the substitutions used to retrieve the username from the certificate field:
 - Username
 - Full username
 - Extended key usage OIDs
 - Use the **Configuration > Authentication > Certificate > Authorization** tab to:
 - Specify the LDAP realm to search
 - Select the **Determine by search** radio button and specify a search filter to map the username to a specific LDAP attribute, such as `(email=$(cs-username))`

- Use the **Configuration > Authentication > Certificate > General tab** to set:
 - Refresh times
 - Inactivity timeout
 - Cookies
 - Virtual URL

Behavior

- ❑ The ProxySG retrieves the end-user PKI certificate from the browser when an HTTP request is received for the domain.
- ❑ The user enters the smart card and pin code information into the browser.
- ❑ The browser retrieves the certificate from a smart card or from within a web browser's certificate store and sends it to the ProxySG.
 - For a specific destination, the certificate must be a valid certificate from a specific Certificate Authority and the certificate must not be revoked.
 - The e-mail address being used as the username must be retrieved from the certificate as a unique ID for the user.
- ❑ The ProxySG does an **LDAP search operation** with the retrieved username from the certificate. If only one entry in the LDAP server exists with this e-mail address, the user is authenticated. If the user has the correct group attributes, the user is authorized to access the Web site.

Chapter 50: Oracle COREid Authentication

This section describes how to configure the ProxySG to consult an Oracle COREid (formerly known as Oracle NetPoint) Access Server for authentication and session management decisions. It contains the following topics:

- ❑ ["About COREid Interaction with Blue Coat" on page 971](#)
- ❑ ["Configuring the COREid Access System" on page 972](#)
- ❑ ["Configuring the ProxySG Realm" on page 973](#)
- ❑ ["Participating in a Single Sign-On \(SSO\) Scheme" on page 973](#)
- ❑ ["Creating a COREid Realm" on page 974](#)
- ❑ ["Configuring Agents for COREid Authentication" on page 975](#)
- ❑ ["Configuring the COREid Access Server" on page 976](#)
- ❑ ["Configuring the General COREid Settings" on page 977](#)
- ❑ ["Creating the CPL" on page 979](#)

About COREid Interaction with Blue Coat

Access to the COREid Access System occurs through the Blue Coat Authentication and Authorization Agent (BCAAA).

Within the COREid Access System, BCAA acts as a custom AccessGate. It communicates with the COREid Access Servers to authenticate the user and to obtain a COREid session token, authorization actions, and group membership information.

HTTP header variables and cookies specified as authorization actions are returned to BCAA and forwarded to the ProxySG. They can (as an option) be included in requests forwarded by the appliance.

Within the ProxySG system, BCAA acts as its agent to communicate with the COREid Access Servers. The ProxySG provides the user information to be validated to BCAA, and receives the session token and other information from BCAA.

Each ProxySG COREid realm used causes the creation of a BCAA process on the Windows host computer running BCAA. When a process is created, a temporary working directory containing the Oracle COREid files needed for configuration is created for that process. A single host computer can support multiple ProxySG realms (from the same or different ProxySGs); the number depends on the capacity of the BCAA host computer and the amount of activity in the realms.

CoreID authentication realms cannot be used to authenticate administrative users to the ProxySG appliance management console.

Note: Refer to the *BCAAA Service Requirements* document for up-to-date information on BCAA compatibility. The *BCAAA Service Requirements* document is posted at MySymantec.

The ProxySG supports authentication with Oracle COREid v6.5 and v7.0.

Configuring the COREid Access System

Note: Blue Coat assumes you are familiar with the configuration of the COREid Access System and WebGates.

Because BCAA is an AccessGate in the COREid Access System, it must be configured in the Access System just like any other AccessGate. BCAA obtains its configuration from the ProxySG so configuration of BCAA on the host computer is not required. If the Cert Transport Security Mode is used by the Access System, then the certificate files for the BCAA AccessGate must reside on BCAA's host computer.

COREid protects resources identified by URLs in policy domains. A ProxySG COREid realm is associated with a single protected resource. This could be an already existing resource in the Access System, (typical for a reverse proxy arrangement) or it could be a resource created specifically to protect access to ProxySG services (typical for a forward proxy).

Important: The request URL is not sent to the Access System as the requested resource; the requested resource is the entire ProxySG realm. Access control of individual URLs is done on the ProxySG using policy.

The COREid policy domain that controls the protected resource must use one of the challenge methods supported by the ProxySG.

Supported challenge methods are Basic, X.509 Certificates and Forms. Acquiring the credentials over SSL is supported as well as challenge redirects to another server.

The ProxySG requires information about the authenticated user to be returned as COREid authorization actions for the associated protected resource. Since authentication actions are not returned when a session token is simply validated, the actions must be authorization and not authentication actions.

The following authorization actions should be set for all three authorization types (Success, Failure, and Inconclusive):

- ❑ A HeaderVar action with the name `BCSI_USERNAME` and with the value corresponding to the simple username of the authenticated user. For example, with an LDAP directory this might be the value of the `cn` attribute or the `uid` attribute.
- ❑ A HeaderVar action with the name `BCSI_GROUPS` and the value corresponding to the list of groups to which the authenticated user belongs. For example, with an LDAP directory this might be the value of the `memberOf` attribute.

After the COREid AccessGate, authentication scheme, policy domain, rules, and actions have been defined, the ProxySG can be configured.

Note: The ProxySG credential cache only caches the user's authentication information for the lesser of the two values of the time-to-live (TTL) configured on the ProxySG and the session TTL configured in the Access System for the AccessGate.

Configuring the ProxySG Realm

The ProxySG realm must be configured so that it can:

- ❑ Communicate with the Blue Coat agent(s) that act on its behalf (hostname or IP address, port, SSL options, and the like).
- ❑ Provide BCAAA with the information necessary to allow it to identify itself as an AccessGate (AccessGate id, shared secret).
- ❑ Provide BCAAA with the information that allows it to contact the primary COREid Access Server (IP address, port, connection information).
- ❑ Provide BCAAA with the information that it needs to do authentication and collect authorization information (protected resource name), and general options (off-box redirection).

For more information on configuring the ProxySG COREid realm, see ["Creating a COREid Realm"](#) on page 974.

Note: All ProxySG and agent configuration occurs on the appliance. The appliance sends the necessary information to BCAAA when it establishes communication.

Participating in a Single Sign-On (SSO) Scheme

The ProxySG can participate in SSO using the encrypted `ObSSOCookie` cookie. This cookie is set in the browser by the first system in the domain that authenticates the user; other systems in the domain obtain authentication information from the cookie and so do not have to challenge the user for credentials. The ProxySG sets the `ObSSOCookie` cookie if it is the first system to authenticate a user, and authenticates the user based on the cookie if the cookie is present.

Since the SSO information is carried in a cookie, the ProxySG must be in the same cookie domain as the servers participating in SSO. This imposes restrictions on the `authenticate.mode()` used on the ProxySG.

- ❑ A reverse proxy can use any `origin` mode.
- ❑ A forward proxy must use one of the `origin-redirect` modes (such as `origin-cookie-redirect`). When using `origin-*-redirect` modes, the virtual URL's hostname must be in the same cookie domain as the other systems. It cannot be an IP address; the default `www.cfauth.com` does not work either.

When using `origin-*-redirect`, the `sso` cookie is automatically set in an appropriate response after the ProxySG authenticates the user. When using `origin` mode (in a reverse proxy), setting this cookie must be explicitly specified by the administrator using the policy substitution variable `$(x-agent-sso-cookie)`. The variable `$(x-agent-sso-cookie)` expands to the appropriate value of the `set-cookie:` header.

Avoiding ProxySG Challenges

In some COREid deployments all credential challenges are issued by a central authentication service. Protected services do not challenge and process request credentials; instead, they work entirely with the `sso` token. If the request does not include an `sso` token, or if the `sso` token is not acceptable, the request is redirected to the central service, where authentication occurs. Once authentication is complete, the request is redirected to the original resource with a response that sets the `sso` token.


If the COREid authentication scheme is configured to use a forms-based authentication, the ProxySG redirects authentication requests to the form URL automatically. If the authentication scheme is not using forms authentication but has specified a challenge redirect URL, the ProxySG only redirects the request to the central service if `always-redirect-offbox` is enabled for the realm on the ProxySG. If the `always-redirect-offbox` option is enabled, the authentication scheme must use forms authentication or have a challenge redirect URL specified.

Note: The ProxySG must not attempt to authenticate a request for the off-box authentication URL. If necessary, `authenticate(no)` can be used in policy to prevent this.

Creating a COREid Realm

To create a COREid realm:

1. Select the **Configuration > Authentication > Oracle COREid > COREid Realms** tab.
2. Click **New**.



3. In the **Realm name** field, enter a realm name. The name can be 32 characters long and composed of alphanumeric characters and underscores. The name *must* start with a letter. The name should be meaningful to you, but it does not have to be the name of the COREid AccessGate.


4. Click **OK** to close the dialog.
5. Click **Apply**.

Configuring Agents for COREid Authentication

You must configure the COREid realm so that it can find the Blue Coat Authentication and Authorization Agent (BCAAA).

To configure the BCAA agent:

1. Select the **Configuration > Authentication > Oracle COREid > Agents** tab.



2. From the **Realm Name** drop-down list, select the COREid realm.
3. Configure the **Primary Agent**:
 - a. In the **Primary agent** section, enter the hostname or IP address where the agent resides.
 - b. Change the port from the default of 16101 if necessary.
 - c. Enter the AccessGate ID in the **AccessGate id** field. The AccessGate ID is the ID of the AccessGate as configured in the Access System.
 - d. If an AccessGate password has been configured in the Access System, you must specify the password on the ProxySG. Click **Change Secret** and enter the password. The passwords can be up to 64 characters long and are always case sensitive.
4. (Optional) Enter an alternate agent host and AccessGate ID in the **Alternate agent** section.

5. (Optional) Select **Enable SSL** to enable SSL between the ProxySG and the BCAAA agent. Select the SSL device profile that this realm uses to make an SSL connection to a remote system. Select any device profile that displays in the drop-down list. For information on using device profiles, see ["About SSL Device Profiles"](#) on page 1291.
6. Specify the length of time in the **Timeout Request** field, in seconds, to elapse before timeout if a response from BCAAA is not received. (The default request timeout is **60** seconds.)
7. If you want username and group comparisons on the ProxySG to be case sensitive, select **Case sensitive**.
8. Click **Apply**.

Configuring the COREid Access Server

After you create a COREid realm, use the COREid Access Server page to specify the primary Access Server information.

To configure the COREid Access Server:

1. Select the **Configuration > Authentication > Oracle COREid > COREid Access Server** tab.

The screenshot shows the 'COREid Access Server' configuration page. The 'REALM name' dropdown is set to 'COREid1'. The 'Protected resource name' field is empty. The 'Security mode' is set to 'Open', and the 'Change Transport Pass Phrase' button is visible. The 'Transport certificates path' field is empty. The 'Always redirect off-box' checkbox is unchecked, and the 'Validate client IP address' checkbox is checked. The 'Add header responses from server' checkbox is unchecked. The 'Access Server ID' field is empty. The 'Access Server hostname' field is empty, and the 'Port' is set to '6021'.

2. Select the realm name to edit from the drop-down list.
3. Enter the protected resource name. The protected resource name is the same as the resource name defined in the Access System policy domain.
4. Select the Security Transport Mode for the AccessGate to use when communicating with the Access System.
5. If Simple or Cert mode is used, specify the Transport Pass Phrase configured in the Access System. Click **Change Transport Pass Phrase** to set the pass phrase.
6. If Cert mode is used, specify the location on the BCAAA host machine where the key, server and CA chain certificates reside. The certificate files must be named `aaa_key.pem`, `aaa_cert.pem`, and `aaa_chain.pem`, respectively.

7. To force authentication challenges to always be redirected to an off-box URL, select **Always redirect off-box**.
8. To enable validation of the client IP address in SSO cookies, select **Validate client IP address**. If the client IP address in the `SSO` cookie can be valid yet different from the current request client IP address because of downstream proxies or other devices, then deselect the **Validate client IP address** in the realm. Also modify the WebGates participating in SSO with the ProxySG. Modify the `WebGateStatic.lst` file to either set the **ipvalidation** parameter to false or to add the downstream proxy/device to the `IPValidationExceptions` lists.
9. If your Web applications need information from the Authorization Actions, select **Add Header Responses**. Authorization actions from the policy domain obtained during authentication are added to each request forwarded by the ProxySG. Header responses replace any existing header of the same name; if no such header exists, the header is added. Cookie responses replace a cookie header with the same cookie name, if no such cookie header exists, one is added.
10. Specify the ID of the AccessGate's primary Access Server.
11. Specify the hostname of the AccessGate's primary Access Server.
12. Specify the port of the AccessGate's primary Access Server.
13. Click **Apply**.

Configuring the General COREid Settings

The COREid General tab allows you to specify a display name, the refresh times, an inactivity timeout value, cookies, and a virtual URL.

To configure the general COREid settings:

1. Select the **Authentication > Oracle COREid > COREid General** tab.

COREid Realms	Agents	COREid Access Server	COREid General
Realm name: COREid1			
Display name: COREid1			
Refresh Times: ☒ Use the same refresh time for all			
Credential refresh time: 900 seconds			
Surrogate refresh time: 900 seconds			
Inactivity timeout: 900 seconds			
Rejected credentials time: 1 seconds			
Cookies			
☐ Use persistent cookies			
☒ Verify the IP address in the cookie			
Virtual URL: www.cfauth.com/			
☒ Challenge user after logout			

- From the **Realm name** drop-down list, select the COREid realm for which you want to change properties.
- If needed, change the COREid realm display name. The default value for the display name is the realm name. The display name cannot be greater than 128 characters and it cannot be null.
- Select the **Use the same refresh time for all** option to use the same refresh time for all.
- Enter the number of seconds in the **Credential refresh time** field. The Credential Refresh Time is the amount of time basic credentials (username and password) are kept on the ProxySG. This feature allows the ProxySG to reduce the load on the authentication server and enables credential spoofing. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, the ProxySG authenticates the user supplied credentials against the cached credentials. If the credentials received do not match the cached credentials, they are forwarded to the authentication server in case the user password changed. After the refresh time expires, the credentials are forwarded to the authentication server for verification.
- Enter the number of seconds in the **Surrogate refresh time** field. The Surrogate Refresh Time allows you to set a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the ProxySG authenticates the transaction. After the refresh time expires, the ProxySG will verify the user's credentials. Depending upon the authentication mode and the user-agent, this may result in challenging the end user for credentials.

The main goal of this feature is to verify that the user-agent still has the appropriate credentials.

7. Type the number of seconds in the **Inactivity timeout** field to specify the amount of time a session can be inactive before being logged out.
8. If you use Basic credentials and want to cache failed authentication attempts (to reduce the load on the authentication service), enter the number of seconds in the **Rejected Credentials time** field. This setting, enabled by default and set to one second, allows failed authentication attempts to be automatically rejected for up to 10 seconds. Any Basic credentials that match a failed result before its cache time expires are rejected without consulting the back-end authentication service. The original failed authentication result is returned for the new request.

All failed authentication attempts can be cached: Bad password, expired account, disabled account, old password, server down.

To disable caching for failed authentication attempts, set the **Rejected Credentials time** field to 0.

9. Select the **Use persistent cookies** check box to use persistent browser cookies instead of session browser cookies.
10. Select the **Verify the IP address in the cookie** check box if you would like the cookies surrogate credentials to only be accepted for the IP address that the cookie was authenticated. Disabling this will allow cookies to be accepted from other IP addresses.
11. Specify the virtual URL to redirect the user to when they need to be challenged by the ProxySG. If the appliance is participating in SSO, the virtual hostname must be in the same cookie domain as the other servers participating in the SSO. It cannot be an IP address or the default, `www.cfauth.com`.
12. Select the **Challenge user after logout** option if the realm requires the users to enter their credentials after they have logged out.
13. Click **Apply**.

Creating the CPL

You can create CPL policies now that you have completed COREid realm configuration. Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes.

The examples below assume the default policy condition is *allow*. On new SGOS 5.x or later systems running the Proxy Edition, the default policy condition is *deny*.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file <Proxy> and other layers.

- ❑ Every COREid-authenticated user is allowed access the ProxySG.

```
<Proxy>
  authenticate(COREidRealm)
```

- ❑ Group membership is the determining factor in granting access to the ProxySG.

```
<Proxy>
  authenticate(COREidRealm)
<Proxy>
  group="cn=proxyusers, ou=groups, o=myco"
  deny
```

Chapter 51: SAML Authentication

SAML 2.0 was developed by the OASIS Security Services Technical Committee. It is an industry standard for retrieving authorization and identity information in XML documents to facilitate single sign-on (SSO) applications or services on the Internet. In SAML authentication, the exchange of information is performed by the following entities:

- ❑ Identity providers (IDPs), which are identity stores. For example, an IDP may have a back-end directory of users. The IDP authenticates the users. Supported IDPs are listed in ["Requirements for SAML Authentication"](#) on page 983.
- ❑ Service providers (SPs), which provide access to applications or services to users. It is the entity that creates an authenticated session for the user. In SGOS 6.5, the ProxySG appliance is the SP.

With SAML authentication, the ProxySG appliance acts as the service provider. SAML realms are not compatible with administrative authentication to the ProxySG appliance management console.

Note: This document assumes that you are familiar with SAML concepts and practices.

The following sections describe how to configure SAML:

- ❑ ["About SAML"](#) on page 982
- ❑ ["Requirements for SAML Authentication"](#) on page 983
- ❑ ["An Overview of the Authentication Process"](#) on page 984
- ❑ ["Export the IDP Metadata File"](#) on page 986
- ❑ ["Prepare the Appliance"](#) on page 988
- ❑ ["Create the SAML Realm"](#) on page 990
- ❑ ["Configure SAML Authorization"](#) on page 992
- ❑ ["Configure the IDP"](#) on page 993
- ❑ ["Prevent Dropped Connections When Policy is Set to Deny"](#) on page 1004
- ❑ ["Backing Up ProxySG Configuration: Considerations for SAML"](#) on page 1004

About SAML

The following sections provide conceptual information you must understand before configuring SAML:

- ❑ ["Federation and Metadata"](#) on page 982
- ❑ ["Assertions"](#) on page 982
- ❑ ["Profiles and Bindings"](#) on page 983

Federation and Metadata

The entities (IDP and SP) must *federate* before authentication can occur. During federation, configuration data is exchanged in *metadata* files. Each entity publishes information about itself in these files and publishes them to a specific location, for example, on the internet or a network drive. When the entities share metadata, they establish and agree on the parameters that they will use for authentication requests and responses. They also share information such as:

- ❑ Entity IDs, which entities use to identify themselves to each other. For example, the Entity ID tells the IDP if an authentication request comes from a federated relying party.
- ❑ The SSO POST endpoint and SSO redirect endpoint to which entities send assertions during authentication. (See ["Assertions"](#) on page 982 for more information.)
- ❑ Each entity's public key certificate, which is used to validate assertions.
- ❑ Whether each entity requires encryption. If one of the entities requires encrypted assertions, it will publish a separate encryption certificate.

Note: The ProxySG appliance may consume encrypted assertions from the IDP, but it does not encrypt authentication requests that it sends to the IDP.

Assertions

The ProxySG appliance and the IDP exchange data in XML documents called *assertions*. After a user is authenticated, the IDP sends an authentication assertion and the ProxySG appliance establishes an authenticated session with the appropriate authorization for the user.

The ProxySG appliance processes SAML authentication responses from the IDP; these responses may contain assertion attributes that describe the authenticated user. For example, `<saml:Attribute Name="mail">` is an assertion attribute that contains the user's email address inside the `<saml:AttributeValue>` element.

You can configure the ProxySG appliance to use assertion attributes in authorization decisions. For more information on attributes, see ["Configure SAML Authorization"](#) on page 992.

Profiles and Bindings

A *profile* contains information about how SAML supports a defined use case. For example, the Web Browser SSO Profile enables single sign-on authentication for resources on the internet.

SAML 2.0 includes protocol-specific *bindings*, which describe how SAML data is exchanged over those protocols. For SGOS 6.5, SAML authentication supports the following for the Web Browser SSO Profile:

- ❑ The HTTP POST binding for authentication responses
- ❑ The HTTP POST binding and the redirect binding for authentication requests

Requirements for SAML Authentication

Setting up a SAML realm for the ProxySG appliance requires the following:

- ❑ One of the following IDPs:
 - Microsoft® Active Directory Federation Services (AD FS) 2.0
Note: ADFS 1.0 ships with Windows Server 2008. If you want to use the SAML realm with AD FS, you must download AD FS 2.0 from the Microsoft website and install it.
 - CA SiteMinder® Federation Partnership R12
 - Oracle® Identity Federation 11g
 - (Available in SGOS 6.5.2 and later) Shibboleth 2.3.5

Checklist: Preparing the IDP

Before you set up a SAML realm, make sure that you have done the following for your IDP:

- ❑ Installed and configured the administration software
- ❑ Set up the identity store for authentication
- ❑ Identified the default user attribute to be passed in SAML assertions, for example, the User Principal Name attribute in LDAP
- ❑ Identified any additional attributes that you want to be passed in assertions, for example, the memberOf attribute, which identifies the groups of which a user is a direct member in LDAP
- ❑ Determined the location (URL) of the IDP's metadata file. This is needed to complete the steps in ["Create the SAML Realm"](#) on page 990. If you import metadata, the realm uses its preconfigured settings. See ["Export the IDP Metadata File"](#) for instructions on locating the metadata file for the IDP.


Note: To import SiteMinder and Oracle metadata, use the `#(config saml <realm-name>) inline idp-metadata <XML>` CLI command. To avoid errors, Symantec recommends that you import metadata through the CLI instead of entering the information manually in the Management Console.

An Overview of the Authentication Process

After you have defined and configured a SAML realm, and both entities have federated and exchanged metadata, authentication can occur.


The following is an overview of what happens when a user goes to a website that requires authentication.

Step 1 - Initial user request



The ProxySG appliance intercepts the user's request and redirects the web browser to the IDP. The redirect URL includes the SAML authentication request that should be submitted to the IDP's SSO service. If the Disable Client Redirect check box is checked, the SG does not redirect the client to the IDP.

Step 2 - Authentication request and response




The IDP asks the SG for the user's credentials, for example by asking for valid login credentials or checking for valid session cookies for stored credentials.

If the appliance responds with valid credentials, the IDP:

- ❑ Signs an authentication response with its private signing key. If the IDP has been configured to send encrypted assertions, the IDP encrypts the assertion before sending it to the ProxySG appliance.

- Sends the authentication response to the ProxySG appliance, which contains the user's username (however, the appliance is not aware of the user's credentials).


Step 3 - Assertion decryption and validation



If the assertion is encrypted, the ProxySG appliance decrypts it. The appliance rejects any unsigned assertions.

The ProxySG appliance validates the assertion and then retrieves the user's name and group memberships (as specified in assertion attributes) from the assertion using the private key.

Step 4 - User request validation



The appliance validates the user request using the corresponding public key, which is embedded in the IDP's signing certificate. Then, the appliance redirects the user to the website and creates an authenticated session for the user.

Set up SAML Authentication

Perform the following steps to set up a SAML realm.

Table 51–1 Steps for Setting up a SAML Realm

Task #	Task	Reference(s)
1	Save a local copy of the IDP's configuration data.	"Export the IDP Metadata File" on page 986
2	Prepare the appliance for SAML authentication: ☐ Import certificates to the CCL. ☐ Set up an HTTPS reverse proxy service. Note: This is required only if the SAML realm is using an HTTPS POST endpoint. ☐ Define assertion attributes. ☐ Create the SAML realm.	• "Export the IDP Metadata File" on page 986 • "Create an HTTPS Reverse Proxy Service" on page 988 • "Configure SAML Attributes" on page 989 • "Create the SAML Realm" on page 990
3	(Optional) To authorize users through one or more SAML realms, specify the criteria to use when searching for users.	"Configure SAML Authorization" on page 992
4	Configure your specific IDP. Note: Shibboleth is supported in SGOS 6.5.2 and later.	• "Configure AD FS" on page 993 • "Configure SiteMinder" on page 996 • "Beside the partnership, select Action > Activate." on page 999 • "Configure Shibboleth" on page 1003
5	Include the SAML realm in your policy.	"Add the SAML Realm to Policy" on page 1004

Export the IDP Metadata File

To export the IDP metadata file, log in to the IDP's administration software.

Exporting IDP metadata entails saving the XML document to disk. It is important to save the metadata file *without* opening it in a browser first. Browsers do not necessarily support XML file structure and may change the XML tags.

If you use SiteMinder, Oracle, or Shibboleth, you will need to copy and paste the metadata file contents to the CLI using the `inline idp-metadata` command. Because XML files are text-based, it is best to use a text editor such as Notepad to open the file to copy its contents.

Note: To ensure that the SAML realm is configured correctly, Symantec recommends that you import metadata instead of entering the information manually. If there are issues with realm configuration, the Authentication debug log shows the following error: **The SAML realm configuration is invalid.**

Export Metadata from AD FS

To export metadata from AD FS:

1. Log in to the AD FS MMC.
2. Select **Endpoints** and look under **Metadata** for the URL beside the **Federation Metadata** type.
3. Copy the URL and paste it into a browser address bar.
4. Save the XML file to a location that the appliance can access.

Export Metadata from SiteMinder

Before you can export metadata, make sure that you have created a SAML 2.0 IDP. This assumes that you have already created the IDP (entity) in SiteMinder.

To export metadata from SiteMinder:

1. Log in to the CA Federation Manager.
2. Select **Federation > Entities**.
3. Beside the entity you created, select **Action > Export Metadata**.
4. In the **Partnership Name** field, enter a name to identify the partnership between the ProxySG and SiteMinder. You will refer to this partnership name later, when you configure the partnership in SiteMinder.
5. Click **Export**. SiteMinder generates the metadata document.
6. Save the XML file to a location that the appliance can access.

Export Metadata from Oracle

To export metadata from Oracle:

1. Log in to the Oracle Enterprise Manager.
2. In the navigation tree on the left, select **Identity and Access > OIF**.
3. On the main page, select **Oracle Identity Federation > Administration > Security and Trust**.
4. Click the **Provider Metadata** tab.
5. In the Generate Metadata section, select **Identity Provider** from the Provider Type menu.
6. Select **SAML 2.0** from the Protocol menu.
7. Click **Generate**. OIF generates the metadata document.
8. Save the XML file to a location that the appliance can access.

Export Metadata from Shibboleth

(Available in SGOS 6.5.2 and later) To export data from Shibboleth:

1. On the server where Shibboleth is installed, browse to the `<shibboleth>/metadata` folder.
2. Copy the `idp-metadata.xml` file.
3. Save the XML file to a location that the appliance can access.

Prepare the Appliance

Complete the following steps to prepare the ProxySG appliance for SAML authentication.

- ❑ ["Configure the CCL" on page 988](#)
- ❑ ["Create an HTTPS Reverse Proxy Service" on page 988](#)
- ❑ ["Configure SAML Attributes" on page 989](#)
- ❑ ["Configure General Settings for SAML" on page 989](#)

Configure the CCL

The ProxySG appliance CCL must contain at least a root certification authority (CA) certificate, but depending on other considerations, you may require more certificates. Refer to the following list to determine which certificates you must import to the CCL.

- ❑ Root CA certificate—**Required**. Add the certificate for the root CA that issued the IDP's signing certificate to the CCL.
- ❑ IDP's signing certificate—**Required if self-signed**. If the IDP's signing certificate is self-signed, add it to the CCL. Certificates signed by the CA are included in SAML assertions.
- ❑ Intermediate CA certificate—**Optional**. You must import intermediate CA certificates to the ProxySG, but it is not necessary to add them to the CCL. For instructions on importing certificates to the ProxySG appliance, see ["Import Certificates onto the ProxySG Appliance" on page 1125](#).

Note: In explicit deployments, if you do not add the certificate for the CA that issued the IDP's certificate to the ProxySG appliance's CCL, HTTPS connections to the IDP fail.

Create an HTTPS Reverse Proxy Service

You should create an HTTPS reverse proxy service only if the SAML realm uses an HTTPS, rather than HTTP, POST endpoint.

Determine whether to use an HTTP or HTTPS POST endpoint:

- ❑ If you use AD FS, the POST endpoint must use HTTPS.
If you use SiteMinder or Oracle, the POST endpoint can use either HTTPS or HTTP.

- ❑ If you want greater security, use an HTTPS POST endpoint.

Note: Regardless of which IDP you use, Symantec recommends using an HTTPS POST endpoint.

Create an HTTPS reverse proxy service to act as the SAML realm's HTTPS POST endpoint. The IDP redirects browsers to this service when it creates assertions.

SSL connections require a certificate and a private key. Browsers must trust the certificate that the HTTPS reverse proxy service uses, and the certificate's **Subject** value must match the **Virtual host** configured in the SAML realm (see "[Create the SAML Realm](#)" on page 990). If the names do not match, SSL hostname mismatch errors occur.

To create the HTTPS reverse proxy service, see "[Creating an HTTPS Reverse Proxy Service](#)" on page 321.

Configure SAML Attributes

The ProxySG appliance maps policy conditions to assertion attribute values. If you require more attributes than the ones included in SAML assertions, you can define them in the SAML realm.

To define assertion attributes:

1. In the Management Console, select **Configuration > Authentication > SAML > Attributes**.
2. Click **New**. The Add SAML Attribute dialog displays.
3. Enter attribute settings:
 - **Attribute name**—This is the name of the attribute as it appears in the ProxySG appliance and IDP configuration, and when referring to the attribute in the `attribute.<name>=` policy condition. The name must be unique.
 - **Attribute data type**—Select **case-exact-string** or **case-ignore-string**. The ProxySG appliance uses this setting to match assertion attribute values with policy conditions.
 - **SAML name**—This is the name of the attribute as it will appear in assertions from the IDP, in the `Name=` XML attribute of the `<Attribute>` element. For example, an assertion might include the line `<saml:Attribute Name="mail">` where `mail` is the SAML attribute name.
4. Click **OK**, and then click **Apply**.

Configure General Settings for SAML

To configure general SAML settings:

1. In the Management Console, select **Configuration > Authentication > SAML > SAML General**.
2. From the **Realm name** menu, select a SAML realm.

3. Configure the following as required.
 - **Display name**—Set the display name of the realm.
 - **Refresh Times**—Do one of the following.
 - Mark the **Use the same refresh time for all** check box to set the same refresh time for credentials and surrogates.
 - Enter different refresh times (in seconds) for credentials and surrogates.
 - **Inactivity timeout**—Enter the number of seconds a session can be inactive before it times out.
 - **Rejected credentials time**—Enter a refresh time (in seconds) for rejected credentials.
 - **Cookies**—Do one or both of the following:
 - **Use persistent cookies**—Mark this to use persistent cookies; leave this unmarked to use session cookies.
 - **Verify the IP address in the cookie**—Mark this to enable verification of cookies' IP addresses.
 - **Challenge user after logout**—Mark this to enable challenging after logout. For example, if this setting is enabled and a user logs out of a web site, the user must enter credentials again the next time they access the web site.
4. Click **Apply** to save your changes to general SAML settings.

Create the SAML Realm

To create the SAML realm:

1. In the Management Console, select **Configuration > Authentication > SAML**.
2. Click **New**.
3. On the Add SAML Realm dialog, in the **Realm name** field, enter a name for the realm. The realm name identifies the realm in areas of the Management Console such as logs and the Visual Policy Manager.
4. Select the trusted CCL. From the **Federated IDP CCL** menu, select the CCL you created in "[Export the IDP Metadata File](#)" on page 986.
5. Do one of the following to specify configuration parameters:
 - (AD FS only) Use preconfigured settings for the IDP. Copy and paste the URL for the metadata into the **Federated IDP metadata URL** field.
 - (SiteMinder, Shibboleth, and Oracle) Import metadata through the `inline idp-metadata` CLI command.
6. From the **Encryption keyring (optional)** menu, select the keyring to use for decrypting encrypted assertions.
7. (Optional) If you require that assertions from the IDP be encrypted, mark the **Require encryption** check box. If you mark the check box, the ProxySG appliance rejects unencrypted assertions.

Note: As long as the encryption keyring is configured, the ProxySG appliance attempts to decrypt encrypted assertions whether or not the **Require encryption** check box is marked.

8. Specify the hostname for the SAML endpoint; in other words, point to the HTTPS reverse proxy listener you set up. In the **Virtual host** field, enter the host and port in format `<hostname_or_IP_address>:<port_number>`. The hostname must match the name of the SSL certificate for the HTTPS reverse proxy service. See ["Create an HTTPS Reverse Proxy Service"](#) on page 988.
9. (Optional) Define limits for assertions' timestamps. Assertions with timestamps that fall outside of these limits are invalid.
 - Specify an interval before the current time. Assertions stamped before this interval are invalid. In the **Not before** field, specify the number of seconds. The default value is 60.
 - Specify an interval after the current time. Assertions stamped after this interval are invalid. In the **Not after** field, specify the number of seconds. The default value is 60.
10. (If applicable) If you defined your own assertion attributes (["Configure SAML Authorization"](#) on page 992), select them from the following menus:
 - **SAML user attribute**—This is the attribute containing the relative username. If you do not specify the attribute, the ProxySG appliance uses the SAML Name ID value for the username.
 - **SAML fullname attribute**—This is the attribute containing the full username.
 - **SAML group attribute**—This is the name of the group membership attribute. Values of this attribute match the `group=` policy condition.
11. (Optional; available in SGOS 6.5.2 and later) Select an **SSL device profile** to use to communicate with the IDP when the redirect/POST URL uses HTTPS.
12. (Optional; available in SGOS 6.5.2 and later) If the client cannot be redirected to, or communicate directly with, the IDP—for example, if a firewall exists between users and the IDP—select **Disable client redirects** to disable the appliance's ability to redirect the browser to the IDP for authentication.

When you disable client redirects, the appliance handles all communication with the IDP and OCS on behalf of the client.
13. (If you selected **Disable client redirects** in the previous step) Determine if you want to prevent the appliance from adding a `BCSI-SWR-` prefix to IDP cookies; this prefix prevents IDP cookies and OCS cookies from having the same name.

The **Prefix IDP cookies** option is available and enabled by default when you select **Disable client redirects**. Clear the **Prefix IDP cookies** option if you do not want to add the prefix to cookies.
14. Click **OK** to save the realm.

Enter Configuration Parameters Manually

Symantec recommends that you import metadata, but as an alternative, you can enter configuration parameters manually after you save the realm.

To enter configuration parameters manually:

1. In the Management Console, select **Configuration > Authentication > SAML**.
2. Select the SAML realm and click **Edit**.
3. In the dialog that displays, specify the SAML entity ID for **Federated IDP entity ID**.
4. Specify one or both of the following endpoints:
 - **Federated IDP POST URL**—The SSO POST endpoint
 - **Federated IDP Redirect URL**—The SSO redirect endpoint
5. Click **OK**.

Configure SAML Authorization

You can authorize users through one or more SAML realms and specify the criteria to use when searching for users.

To configure authorization settings for SAML:

1. In the Management Console, select **Configuration > Authentication > SAML > Authorization**.
2. From the **Realm name** drop-down list, select the SAML realm for which you want to configure authorization settings.
3. (If applicable) To authorize with the current realm, mark the **Self** check box. If you select **Self**, the Authorization username is set automatically to **Use FQDN**.

Note: If you use LDAP for authorization and **Use FQDN** is selected, ensure that the **SAML fullname attribute** (see "[Create the SAML Realm](#)" on page 990) contains the user's distinguished name. Later, you must also configure the IDP to send the distinguished name in assertions.

4. To authorize with another realm, go to the next step.
5. (If applicable) Select the realm with which to authorize from the Authorization realm name menu. Then, choose one of the following options:
 - **Authorization username**—Enter the username in the field.
 - **Use FQDN**—Use the fully-qualified domain name (FQDN).
 - **Determine by search**—Determine the username by LDAP search. Specify the following:
 - **LDAP search realm name**—Enter the name of the LDAP search realm.
 - **Search filter**—Specify the LDAP search filter.

- **User attribute or FQDN**—Specify either the LDAP attribute name or the FQDN as the username attribute for search results.
 - **Set Users to Ignore**—Add, edit, or remove usernames from the list of users to ignore when determining authorization.
6. Click **Apply** to save your changes to SAML authorization settings.

Policy Conditions

The ProxySG uses existing policy conditions to make authorization decisions for the user. These policy conditions map to assertion attribute values.

❑ **group=**

The `group=` condition maps to the values of the **SAML group attribute** setting specified in the realm.

❑ **attribute.<name>=**

The `attribute.<name>=` condition maps to the values of the **Attribute name** setting specified in the realm.

Configure the IDP

This section comprises procedures for configuring your IDP for SAML. Follow the procedures for your deployment. These procedures assume that you have installed and configured the administration software for your IDP.

- ❑ ["Configure AD FS" on page 993](#)
- ❑ ["Configure SiteMinder" on page 996](#)
- ❑ ["Beside the partnership, select Action > Activate." on page 999](#)
- ❑ ["Configure Shibboleth" on page 1003](#)

Note: The procedures for configuring assertion attributes refer to the attribute's SAML name. The SAML name is the name that you specified for the attribute in ["Configure SAML Attributes" on page 989](#).

Configure AD FS

The following steps comprise the minimum required settings to create trust between the ProxySG and AD FS. For other settings that you may require for your deployment, refer to the AD FS documentation.

Note: To perform the procedures in this section, you must be logged in with administrator credentials on the AD FS server.

Import the ProxySG Certificate to AD FS's Trust List

Before the AD FS server can import metadata from the SAML realm, AD FS has to trust the ProxySG appliance's default certificate. To create trust, add the certificate to AD FS's trust list.

The following procedure describes how to add the certificate through Microsoft Internet Explorer 9.x.

To add the certificate to AD FS's trust list in Internet Explorer:

1. In the browser, select **Tools > Internet Options > Content**.
2. Click **Certificates**, and then click **Import**.
3. When you are prompted to specify a store in which to install the certificate, select **Trusted Root Certification Authorities**.

Note: If you do not select the **Trusted Root Certification Authorities** store, any error messages that occur may be inaccurate or unintuitive.

Import Metadata to AD FS

The following procedure describes how to import ProxySG metadata in AD FS:

1. In the AD FS MCC, select **AD FS 2.0 > Trust Relationships > Relying Party Trusts**.
2. Select **Relying Party Trusts**, right click, and then select **Add Relying Party Trust**.
3. On the wizard that displays, click **Start**.
4. Make sure that **Import data about the relying party published online or on a local network** is selected.

5. In the Federation metadata address (host name or URL) field, enter the following URL:

<https://<IP-address>:8082/saml/metadata/<realm-name>/sp>

In the URL, *<IP-address>* is the address of the ProxySG, and *<realm-name>* is the name of the SAML realm.

6. Click **Next**.

Note: If an error message displays when you click **Next**, ensure that the certificate was imported correctly (see "[Import the ProxySG Certificate to AD FS's Trust List](#)" on page 993), and then verify that the hostname you specified in the URL in step 5 matches the certificate's **Subject** value.

In addition, if AD FS fails to validate the certificate, a generic error message displays; the message does not indicate that the certificate is invalid.

7. Enter a display name for the relying party trust and then click **Next**.
8. To allow access to the ProxySG for all users, select **Permit all users to access this relying party**. Do not select this option if you want to limit access to the ProxySG to authorized users. Then, click **Next**.
9. Review your settings, and then click **Next**.
10. Make sure that **Open the Edit Claim Rules** is selected, and then click **Close**.

AD FS prompts you to edit claim rules. See "[Set up Claim Rules for Assertions](#)" on page 995.

Set up Claim Rules for Assertions

Set up a claim rule to send user attributes in SAML assertions:

1. In the AD FS MMC, select **AD FS 2.0 > Trust Relationships > Relying Party Trust**.
2. Select the relying party trust that you created in ["Import Metadata to AD FS"](#) on page 994, right click, and click **Edit Claim Rules**.
3. On the dialog that displays, click **Add Rule**.
4. On the wizard screen that displays, make sure that **Send LDAP Attributes as Claims** is selected for Claim rule template, and then click **Next**.
5. Configure the rule. You can configure the rule to send any attribute, but this procedure describes the following:
 - ["Send User Identity"](#) on page 995
 - ["Send Distinguished Name"](#) on page 995
 - ["Send Group Membership"](#) on page 996

Send User Identity

The following procedure tells you how to pass the `User Principal Name` attribute in the `SAML Name ID` assertion.

1. Specify the following in the claim rule wizard:
 - a. In the **Claim rule name** field, enter the attribute's SAML name.
 - b. For **Attribute store**, select **Active Directory** for the attribute store.
 - c. For **LDAP Attribute**, select the **User-Principal-Name** attribute.
 - d. For **Outgoing Claim Type**, select **Name ID**.
2. Click **Finish**, and then click **OK**.
3. (If required) To add another claim rule, repeat steps 2 through 4 in ["Set up Claim Rules for Assertions"](#) on page 995.

Send Distinguished Name

If you use LDAP for authorization, you must configure a claim rule to send the distinguished name in assertions.

1. Specify the following in the claim rule wizard:
 - a. In the **Claim rule name** field, enter the attribute's SAML name.
 - b. For **Attribute store**, select **Active Directory** for the attribute store.
 - c. For **LDAP Attribute**, enter **distinguishedname** in lower case. (You can type in the drop-down menu.)

Note: Due to a limitation in AD FS, the attribute name disappears if you enter the name once and then go to next field. When this happens, enter the attribute name again exactly as it is shown above.

- d. For **Outgoing Claim Type**, enter DN.
2. Click **Finish**, and then click **OK**.
3. (If required) To add another claim rule, repeat steps 2 through 4 in "[Set up Claim Rules for Assertions](#)" on page 995.

Send Group Membership

You can set up AD FS to send group memberships in assertions. Create another claim rule to specify the attribute, and then add an attribute in SAML.

1. Specify the following in the claim rule wizard:
 - a. In the **Claim rule name** field, enter the attribute's SAML name.
 - b. For **Attribute store**, select **Active Directory** for the attribute store.
 - c. For **LDAP Attribute**, select the **Token Groups-Unqualified Names** attribute.
 - d. For **Outgoing Claim Type**, select **Group**.
2. Click **Finish**, and then click **OK**.
3. (If required) To add another claim rule, repeat steps 2 through 4 in "[Set up Claim Rules for Assertions](#)" on page 995.

Configure SiteMinder

The following steps comprise the minimum required settings to create a partnership between the ProxySG and SiteMinder. For other settings that you may require for your deployment, refer to the CA SiteMinder documentation.

Note: To perform the procedures in this section, you must be logged in with administrator credentials on the SiteMinder server.

Import Metadata to SiteMinder

To import ProxySG metadata to SiteMinder:

1. Go the following URL to export ProxySG metadata:
<https://<IP-address>:8082/saml/metadata/<realm-name>/sp>
Save the file to disk.
2. In the CA Federation Manager, select **Federation > Entities**.
3. Click **Import Metadata**.
4. Beside the Metadata file field, browse to the ProxySG metadata file you saved in the first step.
5. Make sure that the following are selected:
 - For Import As, select **Remote Entity**.
 - For Operation, select **Create New**.
6. Click **Next**.

7. In the Entity Name field, enter a name for the ProxySG as a service provider.
8. Click **Next**.
(The **Import Certificates** step is skipped if the ProxySG metadata doesn't contain a certificate.)
9. Confirm your settings, and then click **Finish**.

Configure the Partnership

Configure the partnership between SiteMinder and the ProxySG.

1. Select **Federation > Partnerships**.
2. Locate the partnership you created when you exported SiteMinder metadata ("[Export Metadata from SiteMinder](#)" on page 987). Beside the partnership name, select **Action > Edit**.
3. Specify the following:
 - For **Remote SP**, select the remote entity you specified when you imported ProxySG metadata.
 - From **Available Directories**, select the LDAP directories you want to use. On the right, you can move directories up or down to dictate the search order.
4. Click **Next**.
5. (Optional) On the **Federation Users** step, specify search filters or users to exclude from accessing the ProxySG.
6. Click **Next**. The next step is Assertion Configuration; see "[Configure Assertions](#)" on page 997.

Configure Assertions

After you configure the partnership, you are at the **Assertion Configuration** step. To send user attributes in SAML assertions:

1. In the Name ID section, for Name ID Format, select the identity format to be passed in SAML assertions.
2. For Name ID Type, select **User Attribute**.
3. In the Value field, specify the primary user attribute, for example, **sAMAccountName** for Active Directory.
4. In the Assertion Attributes section, click **Add Row**.
5. Configure attributes for the assertion. You send any attribute in assertions, but this procedure describes the following:
 - "[Send User Identity](#)" on page 997
 - "[Send Distinguished Name](#)" on page 998
 - "[Send Group Membership](#)" on page 998

Send User Identity

Specify the following:

1. In the Assertion Attribute field, enter the attribute's SAML name.
2. For Type, select **User Attribute**.
3. For Value, enter the name of the attribute you want to send.
4. Click **Next** or add another attribute.

To add another attribute, click **Add Row** in the Assertion Attributes section.

Send Distinguished Name

If you use LDAP for authorization, you must configure SiteMinder to send the distinguished name in assertions.

1. In the Assertion Attribute field, enter the attribute's SAML name.
2. For Type, select **User Attribute**.
3. For Value, enter **distinguishedName**.
4. Click **Next** or add another attribute.

To add another attribute, click **Add Row** in the Assertion Attributes section.

Send Group Membership

Add the `memberOf` attribute to send group membership information in assertions.

1. In the Assertion Attribute field, enter the attribute's SAML name.
2. For Type, select **User Attribute**.
3. For Value, enter **memberOf**.

Note: If any users are members of multiple groups, change the Value to **FMATTR:memberOf**. If you do not modify the attribute value, SiteMinder incorrectly adds the attribute to the assertion, combining all attribute values in a single XML element.

4. Click **Next** or add another attribute.

To add another attribute, click **Add Row** in the Assertion Attributes section.

Select Authentication Mode

After you configure assertions, you are at the **SSO and SLO** step.

1. In the SSO section, select **HTTP-POST** for the SSO Binding.
2. Click **Next**.

Configure Signing and Encryption

After you select the authentication mode, you are at the **Signature and Encryption** step.

1. In the Signature section, for Signing Private Key Alias, select the key that is used to sign assertions.

Alternatively, or if there is no private key in the certificate data store, you can generate a signing key or import one from a PKCS#12 (private key and certificate) file.

2. For Post Signature Options, select **Sign Assertion**.
3. Make sure that **Require Signed Authentication Requests** is not selected.
4. Click **Next**.
5. Confirm your settings and click **Finish**.

Send Encrypted Assertions

If SiteMinder will be sending encrypted assertions to the ProxySG, create an encryption keyring in the SAML realm. For instructions, see "[Creating a Keyring](#)" on page 1119. Then, export ProxySG metadata. Importing the metadata to SiteMinder imports the certificate from the keyring.

Note: You can select the **Show key pair** option so that you can view and copy the keyring as a backup. Backing up the keyring lets you easily import it again should you need to back up the ProxySG configuration later.

To send encrypted assertions:

1. In the Oracle Enterprise Manager, select **Federation > Partnerships**.
2. Beside the partnership you created, select **Action > Edit**.
3. Select the **Signature and Encryption** step.
4. In the Encryption section, select **Encrypt Assertion**.
5. Beside Encryption Certificate Alias, click **Import** and browse to the certificate you exported.
6. Click **Next**.
7. Click **Finish**.

Activate the Partnership

After you have defined the partnership, you must activate it. If users attempt to authenticate while the partnership is not activated, SiteMinder provides an HTTP 500 error message without an explanation of the cause of the error.

1. In the CA Federation Manager, select **Federation > Partnerships**.
2. Beside the partnership you created, select **Action > Activate**. On the confirmation dialog that displays, click **Yes**.

Note: After you activate a partnership, you cannot edit it unless you deactivate it first (**Action > Deactivate**). After you have edited the partnership, it is in an inactive state until you activate it again.

3. Confirm the settings and then click **Finish**.
4. Beside the partnership, select **Action > Activate**.

Configure Oracle

The following steps comprise the minimum required settings for federation between the ProxySG and Oracle. For other settings that you may require for your deployment, refer to the Oracle documentation.

Note: To perform the procedures in this section, you must be logged in with administrator credentials on the Oracle server.

Import Metadata in Oracle

To import ProxySG metadata to Oracle:

1. Go the following URL to export ProxySG metadata:
<https://<IP-address>:8082/saml/metadata/<realm-name>/sp>
Save the file to disk.
2. In the Oracle Enterprise Manager, select **Oracle Identity Federation > Administration > Federations**.
3. In the table of Trusted Providers, click **Add**.
4. In the Add Trusted Provider dialog, beside Metadata Location, click **Choose File** and browse to the location of the ProxySG metadata file that you saved in the previous step.
5. Click **OK**.

Set up Federation Between the ProxySG Appliance and Oracle

To set up federation, first configure the ProxySG as a service provider. Then, configure Oracle as the identity provider.

1. In the Oracle Enterprise Manager, select **Oracle Identity Federation > Administration > Federations**.
2. In the table of Trusted Providers, select the Provider ID you added when the ProxySG metadata was imported, and then click **Edit**.
3. Select **Update Provider Manually**.
4. For Provider Types, make sure that both **Service Provider** and **Authentication Requester** are selected.
5. On the Oracle Identity Federation Settings tab, select **Enable Attributes in Single Sign-On (SSO)**.
6. In the list of attributes, select **Email Address** and clear all other selections.
7. Click **Apply**.
8. Select **Oracle Identity Federation > Administration > Identity Provider**.
9. On the SAML 2.0 tab, select **Enable Identity Provider**.
10. In the Assertion Settings section, select **Send Signed Assertion**.
11. Click **Apply**.

Set up Attribute Mappings

When you set up attribute mapping, you specify the name with which an attribute should be defined in the SAML assertions.

To set up attribute mappings:

1. In the Oracle Enterprise Manager, select **Oracle Identity Federation > Administration > Federations**.
2. Select the Provider ID for the ProxySG and then click **Edit**.
3. On the Oracle Identity Federation Settings tab, next to Attribute Mappings and Filters, click **Edit**.
4. Configure the assertion. You can send any attribute, but this procedure describes the following:
 - ["Send User Identity"](#) on page 1001
 - ["Send Distinguished Name"](#) on page 1001
 - ["Send Group Membership"](#) on page 1002

Send User Identity

Specify the following:

1. On the Name Mappings tab, click **Add**.
2. On the Add Attribute Name Mapping dialog, specify the following:
 - For User Attribute Name, enter **sAMAccountName**.
 - For Assertion Attribute Name, enter the attribute's SAML name.
 - Select **Send with SSO Assertion** in order for the attribute to appear in the assertion.
3. Click **OK**.
4. Click **OK** to save your changes. Alternatively, click **Add** to add another attribute.

Send Distinguished Name

Specify the following:

1. Click **Add**.
2. On the Add Attribute Name Mapping dialog, specify the following:
 - For User Attribute Name, enter **distinguishedName**.
 - For Assertion Attribute Name, enter the attribute's SAML name.
 - Select **Send with SSO Assertion** in order for the attribute to appear in the assertion.
3. Click **OK**.
4. Click **OK** to save your changes. Alternatively, click **Add** to add another attribute.

Send Group Membership

Add the `memberOf` attribute to send group membership information in assertions:

1. Click **Add**.
2. On the Add Attribute Name Mapping dialog, specify the following:
 - For User Attribute Name, enter **memberOf**.
 - For Assertion Attribute Name, enter the attribute's SAML name.
 - Select **Send with SSO Assertion** in order for the attribute to appear in the assertion.
3. Click **OK**.
4. Click **OK** to save your changes. Alternatively, click **Add** to add another attribute.

Sign Outgoing Assertions

If you have not already set up Oracle with a signing certificate, select a keystore to sign outgoing assertions. If you have already set up a signing certificate, skip this procedure.

1. In the Oracle Enterprise Manager, select **Oracle Identity Federation > Administration > Security and Trust**.
2. Click **Wallet**.
3. Click **Update**.
4. In the Update Wallet dialog, in the Signature section, select a keystore that contains the certificate and private key to use for signing outgoing assertions.

Click **ok**. Send Encrypted Assertions

If Oracle will be sending encrypted assertions to the ProxySG, create an encryption keyring in the SAML realm. For instructions, see ["Creating a Keyring"](#) on page 1119. Then, export ProxySG metadata. Importing the metadata to Oracle imports the certificate from the keyring.

To send encrypted assertions after you have imported the certificate:

1. In the Oracle Enterprise Manager, select **Oracle Identity Federation > Administration > Federations**.
2. In the table of Trusted Providers, select the Provider ID for the ProxySG, and then click **Edit**. On the Oracle Identity Federation Settings tab, scroll down to **Identity Provider/Authority Settings**.
3. Verify the following options.
 - In the Assertion Settings list, select **Send Encrypted Assertions**.
 - In the Protocol Settings list, select **Include Signing Certificate in XML Signatures**.
 - In the Messages to Send/Require Signed list, beside Response with Assertion - HTTP POST, clear **Send Signed**.
4. Click **Apply**.

Note: Be sure to update the metadata after enabling encrypted assertions if the last metadata file you uploaded didn't include the encryption keyring.

Configure Shibboleth

(Available in SGOS 6.5.2 and later) Import Shibboleth's certificate to the ProxySG appliance:

1. Copy the contents of `<shibboleth>/conf/idp.cert`.
2. In the Management Console, select **Configuration > SSL > CA Certificates**. Click **Import**.
3. Click **Paste from Clipboard**. Click **OK**.
4. Include this certificate in a CCL. See "Export the IDP Metadata File" on page 986.

Configure the Partnership

Configure the partnership:

1. Download the ProxySG's SAML realm metadata. It is located in <https://<sg-ip>:8082/saml/metadata/<realm-name>/sp>.
2. In the Management Console, select **Statistics > Advanced > SAML2 > SP metadata for SAML2 realms**. Note that the metadata contains the ID ("entityID" in EntityDescriptor element) and the virtual protocol, hostname, and port number ("Location" attribute in "AssertionConsumerServiceElement") required in the next step.
3. Copy the metadata to `<shibboleth>/metadata/<my-metadata>.xml`.
4. Add a new relying party on Shibboleth. Add the following in `<shibboleth>/conf/relying-party.xml`:

```
<rp:RelyingParty provider="<virtual-host>/saml/<realm-name>"
id="<virtual-host>/saml/<realm-name>"
defaultSigningCredentialRef="IdPCredential">
  <rp:ProfileConfiguration xsi:type="saml:SAML2SSOProfile"
includeAttributeStatement="true"
  assertionLifetime="PT5M" assertionProxyCount="0"
  signResponses="never" signAssertions="always"
  encryptAssertions="never" encryptNameIds="never"/>
</rp:RelyingParty>
```

5. Add the XML element `<metadata:MetadataProvider>` in `<shibboleth>/conf/relying-party.xml`. Note that the "id" attribute must be unique among other existing metadata providers.

```
<metadata:MetadataProvider id="<an id>"
xsi:type="metadata:ResourceBackedMetadataProvider">
  <metadata:MetadataResource xsi:type="resource:FilesystemResource"
file="<shibboleth>/metadata/<my-metadata>.xml"/>
</metadata:MetadataProvider>
```

Add the SAML Realm to Policy

After completing SAML realm configuration, you can install policy using content policy language (CPL). Be aware that the examples below are just part of a comprehensive authentication policy.

Note: The examples below assume that the default policy condition is `allow`.

Refer to the *Content Policy Language Reference* for details about CPL and how transactions trigger the evaluation of policy file layers.

Note: SAML realms do not work with any “form” authentication mode.

To allow every SAML-authenticated user access to the ProxySG:

```
<Proxy>
  authenticate(<realm-name>)
```

To specify group membership as the determining factor in granting access to the ProxySG:

```
<Proxy>
  authenticate(<realm-name>)
<Proxy>
  group="saml_users" ALLOW
  deny
```

In the examples above, `<realm-name>` is the name of the SAML realm you created.

Prevent Dropped Connections When Policy is Set to Deny

Users might experience dropped connections if the ProxySG uses the default policy of **Deny**. When the policy is set to **Deny**, the ProxySG intercepts and denies requests to the IDP.

To prevent dropped connections due to this limitation, install the following policy to allow requests to the IDP:

```
<Proxy>
  authenticate(<realm-name>)
<Proxy>
  allow group=saml_users
<Proxy>
  allow url.host=<hostname>
```

In the policy example above, `<realm-name>` is the name of the SAML realm and `<hostname>` is the hostname of the IDP.

Backing Up ProxySG Configuration: Considerations for SAML

You may need to back up the ProxySG configuration and save the backup file (called an *archive*) on a remote system, which you can restore in the unlikely event of system failure or replacement. For more information on configuration backups, see ["Backing Up the Configuration"](#) on page 69.

Save Keyrings Before Backing up the Configuration

Backing up the configuration does not automatically save keyring data; thus, you must save all keyrings and certificates before creating a configuration archive.

To save keyrings before backup:

1. In the command line interface (CLI), enter configuration mode and issue the following command:

```
#(config ssl)show ssl keypair <keyring_name>
```
2. Copy and paste the output from the command into a text editor. You will copy and paste this text into the CLI after you restore the appliance.

Import Saved Keyrings After Restoring the Configuration

After restoring the ProxySG appliance, but *before* applying the archived configuration, issue the following CLI command:

```
#(config ssl)inline keyring show <keyring_name> <eof marker>
<copy and paste text here>
<eof marker>
```

You should now be able to apply the archived configuration without having to create and import a new keyring.

Re-Import the ProxySG Certificate to the Trust List (AD FS)

In order to establish an HTTPS connection to the ProxySG appliance, AD FS must trust the default certificate set in the ProxySG Management Console; however, after you back up and restore a configuration archive, you must re-import the certificate to AD FS. For instructions, see ["Import the ProxySG Certificate to AD FS's Trust List "](#) on page 993.

Chapter 52: Integrating the Appliance with Your Windows Domain

The following configurations require that you join your ProxySG appliance to join your Windows Domain:

- ❑ To accelerate encrypted MAPI traffic, the ProxySG appliance at the branch office must join the same domain as the Exchange server. For details on all the required steps for accelerating encrypted MAPI, see ["Optimizing Encrypted MAPI Traffic"](#) on page 282.
- ❑ If you want the ProxySG appliance to perform Integrated Windows Domain Authentication (IWA) by directly accessing your Active Directory (AD) rather than using the Blue Coat Authentication and Acceleration Agent (BCAAA), you must first join the appliance to your Windows domain. For more information, see ["Configuring a Direct Connection to the Windows Domain"](#) on page 1022. If you want to authenticate users in different AD domains that do not have trust relationships, you must join the appliance to each domain.

Integrate the ProxySG Appliance into the Windows Domain

To integrate the ProxySG appliance into one or more Windows domains, you must complete the following tasks:

1. ["Synchronize the ProxySG Appliances and DC Clocks"](#) on page 1007
2. ["Join the ProxySG Appliance to the Windows Domain"](#) on page 1008

Synchronize the ProxySG Appliances and DC Clocks

The ProxySG cannot join a Windows domain unless its internal clock is in sync with the Domain Controller. To ensure that the ProxySG clocks are synchronized with the Domain Controller clock, use either of the following techniques:

- ❑ Specify the same NTP servers for the ProxySG appliances and the Domain Controller.
- ❑ Configure the ProxySG appliances to use the Domain Controller as the NTP source server.

The ProxySG NTP configuration options are located on the **Configuration > General > Clock** tab.


Join the ProxySG Appliance to the Windows Domain

After you have synchronized the ProxySG appliance's internal clock with the Domain Controller, you can join the appliance to one or more Windows domains as follows:

1. From the ProxySG Management Console, select **Configuration > Authentication > Windows Domain > Windows Domain**.
2. In the **Hostname** panel, specify the hostname to use:
 - (Recommended) Select **Use Default - {SG-serial_number}** to use the default hostname.
 - Select or specify a different hostname.


Note: Unless you have a specific need to use a particular hostname (for example, to ensure correct DNS lookup), Symantec recommends that you use the default hostname to guarantee that each appliance's hostname is unique. In addition, you must use unique hostnames for multiple ProxySG appliances joined to the same domain.

3. Click **Apply**.
4. Click **Add New Domain**. The Add Windows Domain dialog displays.
5. Enter a **Domain name** alias and then click **OK**.



6. To save the domain alias setting, click **Apply** and then click **OK**. You will not be able to join the domain until you have saved the domain alias setting.


7. Select the domain **Name** you created and click **Join**. The Add Windows Domain dialog displays.



8. Configure the domain membership information:
 - a. In the **DNS Domain Name** field, enter the DNS name for the Windows Active Directory domain. This is *not* the fully qualified domain name of the ProxySG.

Note: The ProxySG appliance must be able to resolve the DNS domain name you supply for the Active Directory domain or the appliance will not be able to join the domain. If DNS resolution fails, check your DNS configuration.

- b. Enter the primary domain access **User Name**. You can either enter the plain user name (for example, administrator) or use the `username@dnsname` format (administrator@acme.com). This account must have rights for joining the domain.
 - c. Enter the **Password** for this user.
 - d. Click **OK**. The appliance displays a message indicating that the domain was successfully joined and the value in the **Joined** field changes to **Yes**.



9. If you want to add additional Windows domains, repeat steps 3 through 8.
10. Click **Apply** to save your changes.

Edit a Windows Domain

The SG administrator can configure a MaxConcurrentApi value for each of the joined Windows domains on the SG. The default value is 2 (same as the default for a Windows member server). The minimum value is 2 and the maximum value is 150.

The SG administrator can also specify a preferred Schannel DC and alternate Schannel DC for each domain. If the preferred Schannel DC is available, the SG will always connect to it, even if it sees another DC that appears to be faster. That serves two purposes:

- The customer need only increase Maximum number of concurrent Schannel connections on the preferred and alternate DCs, rather than on every DC in the domain.
- The preferred and alternate DCs can be read-only. Some customers are willing to deploy a read-only DC that is dedicated just to handling authentication requests for an SG, whereas they would not be willing to deploy a regular “writable” DC for that purpose.

If the SG cannot connect to either of its preferred or alternate DCs, then it will connect to the fastest available DC. The SG will periodically check to see if the preferred or alternate DC comes back online, and will reconnect if it does.

All options in the Edit Windows Domain dialog box are optional.

1. From the ProxySG Management Console, select **Configuration > Authentication > Windows Domain > Windows Domain**.
2. Select a domain in the Domains list and click **Edit**.

Note: Domain controller options are for NTLM authentication only

3. Enter the preferred controller in the **Preferred domain controller** text box.
4. Enter an alternate domain controller in the **Alternate domain controller** text box. The alternate domain controller is used if the preferred domain controller is not available.

The preferred and alternate domain controllers can be read-only. However, if you use a read-only domain controller, you need to replicate user passwords to that domain controller. If the domain controller doesn’t have a copy of the user’s password, it must forward the request to a writable domain controller that has a copy, which will diminish performance. Consult Microsoft documentation to figure out how to do this in your environment.

<http://technet.microsoft.com/en-us/library/cc732801%28v=ws.10%29.aspx>

5. Enter the **Maximum number of concurrent Schannel connections**. The range is 2-150.

Note: In order for the maximum number of concurrent connections to take effect, you must enter the same number in the registry for the Domain Controller(s). The registry setting on the Domain Controller is MaxConcurrentAPI. If you change the MaxConcurrent API setting, you must restart the NetLogon service on the Domain Controller, or reboot the Domain Controller after changing the MaxConcurrent API setting.

6. Click **OK**.

Section 1 Configure SNMP Traps for the Windows Domain

(Introduced in SGOS 6.5.7.1) You can enable SNMP traps for the Windows domain to be notified when errors or issues occur. If SNMP is enabled, you can specify thresholds for any latency and authentication failures that occurred within a given period of time:

- ☐ Last minute
- ☐ Last 3 minutes
- ☐ Last 5 minutes
- ☐ Last 15 minutes
- ☐ Last 60 minutes

Configure SNMP traps:

1. In the Management Console, select **Configuration > Authentication > Windows Domain > SNMP**.
2. Select a domain from the Domain drop-down list.
3. Select **Enable SNMP**.
4. Specify values (in milliseconds) for each time interval for the following:
 - **Average Latency**
 - **Minimum Latency**
 - **Maximum Latency**
 - **Authentication Failures**
5. (Optional) Specify thresholds for secure channel (Schannel):
 - **Schannel Resets Threshold**
 - **Schannel Timeouts Threshold**
 - **Schannel Waiters Threshold**
6. To save your settings, click **Apply** and then click **OK**.

Chapter 53: Integrating ProxySG Authentication with Active Directory Using IWA

Integrated Windows Authentication (IWA) is an authentication scheme that allows you to authenticate and authorize users against your Windows Active Directory (AD). One of the main benefits of IWA is that it can provide a single sign-on experience for your users. When configured properly, the user agent or browser will automatically provide the users' domain credentials to the appliance when challenged without prompting the end users.

Another benefit of IWA is that it provides authorization without any additional configuration because it automatically returns group membership information for the user as part of the authentication response. The ProxySG appliance can then use this group membership information to enforce its authorization policies.

Symantec supports two methods to integrate the appliance with Active Directory using IWA:

- ❑ IWA BCAAA - Connect via an authentication agent running on a Windows server in your domain.

For instructions, refer to the *BCAAA Service Requirements* document posted on MySymantec:

https://support.symantec.com/content/unifiedweb/en_US/Documentation.html?prodRefKey=1145522

- ❑ IWA Direct - Connect the appliance directly to your AD domains on Windows Server 2008, 2012, or 2016. Refer to the following sections.

You can authenticate administrative users to the ProxySG appliance with IWA realms. For maximum security, Symantec recommends using an IWA/BCAAA configuration, where the communication between the appliance and BCAAA is secured with TLS. IWA Direct, though less secure but also compatible with administrative authentication.

The following sections describe how to configure IWA:

- ❑ "About IWA" on page 1014
- ❑ "Preparing for a Kerberos Deployment" on page 1016
- ❑ "Configuring IWA on the ProxySG Appliance" on page 1018
- ❑ "Creating the IWA Authentication and Authorization Policies" on page 1027
- ❑ "Configuring Client Systems for Single Sign-On" on page 1033
- ❑ "Using IWA Direct in an Explicit Kerberos Load Balancing/Failover Scenario" on page 1034

About IWA

The following sections provide the conceptual information you must understand before configuring IWA:

- ❑ ["About IWA Challenge Protocols"](#) on page 1014
- ❑ ["About IWA Failover"](#) on page 1014

About IWA Challenge Protocols

When configured for IWA, the ProxySG appliance determines which of the following protocols to use to obtain Windows domain login credentials each time it receives a client request that requires authentication:

- ❑ **Kerberos**—This is the most secure protocol because it establishes mutual authentication between the client and the server using an encrypted shared key. This protocol requires additional configuration and the appliance will silently downgrade to NTLM if Kerberos is not set up properly or if the client cannot do Kerberos. For more information, see ["Preparing for a Kerberos Deployment"](#) on page 1016.
- ❑ **NTLM**—Uses an encrypted challenge/response that includes a hash of the password. NTLM requires two trips between the workstation and the appliance, and one trip between the appliance and the Domain Controller. It therefore puts more load on the network than Kerberos, which only requires one trip between the workstation and the appliance, and doesn't require a trip between the appliance and the Domain Controller.
- ❑ **Basic**—Prompts the user for a username and password to authenticate the user against the Windows Active Directory.

When the ProxySG appliance receives a request that requires authentication, it consults the IWA configuration settings you have defined to determine what type of challenge to return to the client. It will try to use the strongest authentication protocol that is configured and, if the browser cannot use that protocol or if it is not configured properly, the appliance will downgrade to the next authentication protocol. For example, if you configure the IWA realm to allow Kerberos and NTLM authentication, but the user agent/browser does not support Kerberos, the appliance will automatically downgrade to NTLM.

IWA authentication realms (with basic credentials) can be used to authenticate administrative users (read only and read/write) to the management console. To ensure that credentials are not sent in clear text, configure the IWA realm to use TLS to secure the communication with the BCAAA server, or in the case of IWA direct, secure the communication from the appliance to the domain.

About IWA Failover

The way IWA failover works depends on your deployment:

- ❑ ["IWA Direct Failover"](#) on page 1015
- ❑ ["IWA BCAAA Failover"](#) on page 1015

IWA Direct Failover

For IWA Direct, the realm is considered “healthy” if the ProxySG appliance is able to establish a connection to the Windows domain to which it is a member. As with any other device in the Windows domain, the ProxySG appliance will establish a connection with the closest Windows Domain Controller upon successful domain login. If the Domain Controller to which the ProxySG appliance is connected goes down, the appliance will send an LDAP ping to locate and connect to the next closest Domain Controller.

Because communication between the ProxySG appliance and the Windows Active Directory relies on DNS, you must make sure that the appliance is configured to use more than one DNS server to ensure proper failover. This will ensure that the ProxySG appliance will still be able to communicate with AD, should the primary DNS server go offline. For instructions, see ["Adding DNS Servers to the Primary or Alternate Group"](#) on page 816.

IWA BCAAA Failover

For IWA BCAAA, the realm is considered “healthy” (and therefore won’t fail over) if the ProxySG appliance is able to establish a connection to the BCAAA service. This means that the ProxySG appliance is able to complete the TCP handshake with BCAAA on port 16101 (or whichever port the BCAAA service is configured to use), and the appliance has been able to send BCAAA its “login” message. There are several different failover scenarios in a BCAAA deployment:

- ❑ Each time a ProxySG appliance connects to BCAAA, the BCAAA service (bcaaa.exe) spawns a new BCAAA process (i.e. bcaaa-130.exe). If the BCAAA process crashes, the TCP connection with the corresponding ProxySG appliance will be reset and the ProxySG appliance will attempt to reconnect to the BCAAA service. Other ProxySG appliances that are connected to other instances of the BCAAA process will be unaffected.
- ❑ If the BCAAA service (bcaaa.exe) crashes or is stopped, but the Windows system on which it is running remains available, any ProxySG appliance that is already connected to the BCAAA process (i.e. bcaaa-130.exe) will have their connections reset. The ProxySG appliances will not be able to reconnect to BCAAA because the service is no longer running, and will instead fail over to the secondary BCAAA server.
- ❑ If the Windows server on which BCAAA is running crashes or becomes unavailable, it cannot reset the TCP connection. In this case, BCAAA must wait for the ProxySG appliance’s TCP connection to the Windows server to time out. This can take a couple of minutes, and won’t occur until the ProxySG appliance attempts to send a new authentication request.
- ❑ If the BCAAA server loses its connection to the Windows Domain Controller, it will automatically fail over to a different Domain Controller. Keep in mind that BCAAA cannot detect when Windows fails to connect to any Domain Controllers in a particular domain. In this case all authentication requests will fail, but because the connection between the BCAAA service and the ProxySG appliance is still considered healthy, the ProxySG will not fail over to the secondary BCAAA service.

In addition, authentication requests can be slowed significantly if BCAA is querying a slow Domain Controller. However, this will not cause the ProxySG appliance to fail over to the secondary BCAA server. By default, BCAA will query whichever Domain Controller is chosen at boot time by the server it is installed on, and it only changes if the Domain Controller goes down or the server reboots. You can see and/or modify what Domain Controller the BCAA server is communicating with using the `nltest.exe` utility, which is part of the Windows Support Tools.

To see which Domain Controller the BCAA server is communicating with:

```
nltest /sc_query:internal.domain.com
```

To switch to a different Domain Controller:

```
nltest /sc_reset:internal.domain.com\new_dc_name
```

Preparing for a Kerberos Deployment

Kerberos is the recommended authentication protocol for IWA because it is more secure than NTLM or Basic and it puts the least load on your network.

To ensure that IWA uses the Kerberos protocol rather than downgrading to NTLM, you just need to make sure that authentication requests are directed to the Kerberos service associated with the ProxySG appliance. The way you do this depends on how your IWA realm is connecting to the Active Directory as follows:

- ❑ ["Enabling Kerberos in an IWA Direct Deployment" on page 1016](#)
- ❑ ["Enabling Kerberos in a BCAA Deployment" on page 1017](#)

Enabling Kerberos in an IWA Direct Deployment

In an IWA Direct realm, Kerberos configuration is minimal because the appliance has its own machine account in Active Directory and it uses its account password to decrypt service tickets from clients. Therefore, there is no need for you to create a privileged Active Directory account or generate a service principal name (SPN) for the appliance as is required with an IWA BCAA realm.

To ensure that IWA uses the Kerberos protocol rather than downgrading to NTLM, you just need to make sure that authentication requests are directed to the DNS name of the appliance's Active Directory machine account name as follows:

1. Create a DNS "A" record for the ProxySG that resolves to the DNS name of the appliance's Active Directory machine account name. For example, if you have an appliance named ProxySG1 with IP address 1.2.3.4 in the blue9 Active Directory domain at acme.com, you would create the following DNS record:

ProxySG1.blue9.acme.com	Host (A)	1.2.3.4

2. Ensure that client requests are directed to the DNS name for the ProxySG's Active Directory machine account:
 - **Explicit deployments**—Configure the client browser explicit proxy settings to point to this DNS name.

- **Transparent deployments**—Set the **Virtual URL** in the realm configuration (on the **IWA General** tab) to this DNS name. In addition, make sure that the DNS name for the ProxySG appliance's Active Directory domain is either included in the workstation's list of imputing DNS suffixes or explicitly specified as part of IE's local intranet zone. For example, if your AD domain DNS name is `blue9.acme.com`, then you would add `*.blue9.acme.com` to IE's local intranet zone. See [Step 6](#) on page 1027 in "Defining IWA Realm General Properties".

Enabling Kerberos in a BCAAA Deployment

For the BCAAA service to participate in an IWA Kerberos authentication exchange, it must share a secret with the Kerberos server (called a KDC) and have registered an appropriate Service Principal Name (SPN).

To prepare for a BCAAA Kerberos deployment:

1. Create a DNS "A" record for the ProxySG that resolves to the appliance's Fully Qualified Domain Name (FQDN). Keep in mind that the DNS name you choose must not match the Active Directory machine account name for the appliance. For example, rather than using the machine name, you might create a DNS entry for the appliance using a name such as `bcaaaUser1`. Supposing the appliance is in the `acme.com` domain and has an IP address of `1.2.3.4`, you would create the following DNS record:

<code>bcaaaUser1.acme.com</code>	Host (A)	<code>1.2.3.4</code>

After you create the DNS mapping, make sure you can ping the appliance using the FQDN.

2. Create a domain user account for the BCAAA service in the Windows Active Directory (AD).
3. Install BCAAA. Refer to the BCAAA Service Requirements document for installation instructions. The BCAAA Service Requirements document is posted at MySymantec.
4. Configure the BCAAA Windows service on the system where you just installed BCAAA to log on using the domain user account you created for it in [Step 2](#) rather than using the local system account.
5. In the Local Security Policy of the server on which BCAAA is running, modify the user rights assignment for the BCAAA domain user to have the following rights:
 - Full access to the directory where you installed BCAAA
 - Act as part of the operating system (not required for BCAAA 6.0)
 - Log on as a service
6. Register the Kerberos Service Principal Name (SPN) for the ProxySG appliance:
 - a. Log in to the Domain Controller using an account with administrative access and open a command prompt.
 - b. Enter the following case-sensitive command:

```
setspn -A HTTP/<FQDN_of_ProxySG> <AD_Account_Name>
```

Where <FQDN_of_ProxySG> is the FQDN of the ProxySG appliance as specified in the browser's explicit proxy configuration (explicit deployments) or in the Virtual URL setting in the IWA realm configuration (transparent deployments) and <AD_Account_Name> is the name of the BCAA domain service account.

For example:

```
setspn -A HTTP/bcaaaUser1.acme.com AcmeDomain\BCAAUser
```

Note: Do not assign the same SPN to multiple Active Directory accounts or the browser will fall back to NTLM without providing any warning or explanation. To list all SPNs that are currently registered on an account, use the `setspn -L <AD Account Name>` command. If you find a duplicate, remove the extraneous SPN using the `setspn -D <SPN>` command.

Configuring IWA on the ProxySG Appliance

To set up IWA between the ProxySG appliance and your Active Directory, you must complete the following tasks:

- ❑ ["Creating an IWA Realm" on page 1018](#)
- ❑ ["Configuring IWA Servers" on page 1020](#)
- ❑ ["Defining IWA Realm General Properties" on page 1025](#)

Creating an IWA Realm

Before you can create an IWA realm, you must integrate with the Windows domain. The way you do this depends on how you plan to connect to your Active Directory:


- ❑ **Direct**—The ProxySG appliance will communicate directly with your Domain Controllers to obtain authentication information. Before you can use this option, you must join the ProxySG appliance to the Windows domains that contain your users as described in ["Integrate the ProxySG Appliance into the Windows Domain" on page 1007](#).

Note: Refer to the *BCAAA Service Requirements* document for up-to-date information on BCAA compatibility. The *BCAAA Service Requirements* document is posted at MySymantec.

- ❑ **BCAAA**—The ProxySG appliance will contact the BCAA server when it needs to authenticate a user. To use this option, you must first install BCAA on a dedicated server in your Windows domain and configure it to communicate with both the DC and with the appliance as an authentication agent. Use this option if you do not want to allow the ProxySG appliance to join your Windows domain. Refer to the BCAA Service Requirements document for more information. The *BCAAA Service Requirements* document is posted at MySymantec.

To create an IWA realm:

1. Select **Configuration > Authentication > IWA > IWA Realms**.
2. Click **New**.



3. Enter a **Realm name**. The name can be 32 characters long and composed of alphanumeric characters and underscores. The name *must* start with a letter.
4. Select the type of **Active Directory Connection** you are using and then provide the appropriate configuration information as follows:
 - **Direct**—Select this option if you want the appliance to connect directly to the Windows **Domain** to obtain authentication information. If you have not yet joined the ProxySG appliance to at least one Windows domain, you will not be able to select this option.
 - **BCAAA**—In the **Primary server host** field, enter the hostname or IP address of the server where you installed BCAA. In addition, if you configured BCAA to use a port other than the default (16101), change the value in the **Port** field to match what you configured on BCAA.

Note: If you plan to secure communication between the appliance and BCAA, use a host name rather than an IP address. The DNS server that the appliance is configured to use must be able to resolve the hostname.

5. Click **OK** to close the dialog.
6. To save your settings, click **Apply**.

Configuring IWA Servers

You use the **IWA Servers** tab to configure the connection between the ProxySG appliance and the authentication server (either directly or via BCAA) and to specify the type of credentials to accept from the browser/user agent. You can also verify your configuration from this tab.

The way you set up the configuration depends on whether you connecting directly to the Domain Controller or you are using BCAA to connect to the domain:

- ❑ ["Connecting to the Windows Domain using BCAA" on page 1020](#)
- ❑ ["Configuring a Direct Connection to the Windows Domain" on page 1022](#)

Connecting to the Windows Domain using BCAA

If you plan to use a BCAA server to act as an intermediary between your ProxySG appliance and your Active Directory, you can configure and verify the authentication settings as follows:

1. Select the **Configuration > Authentication > IWA > IWA Servers** tab.

The screenshot shows the 'IWA Servers' configuration page. At the top, there are three tabs: 'IWA Realms', 'IWA Servers' (which is active), and 'IWA General'. Below the tabs, the 'Realm name' is set to 'iwa' and the 'Domain' is set to 'US1'. Under 'Client Authentication Modes', three checkboxes are checked: 'Allow Basic credentials', 'Allow NTLM credentials', and 'Allow Kerberos credentials'. The 'Load-balanced Kerberos' section has a text field for 'Configure ProxySG Kerberos credentials (typically used in load balancing.)' and two buttons: 'Set credentials' and 'Clear credentials'. The 'Test Configuration' section has a button labeled 'Test Configuration'. At the bottom, there is a 'Timeout request after' field set to '60' seconds.

2. From the **Realm name** drop-down list, select the IWA realm you want to configure. If you have not yet created a realm, see ["Creating an IWA Realm" on page 1018](#).
3. If you have not yet installed a primary BCAA server and, optionally, a secondary BCAA server, you must do so before proceeding. Use the **Click here to download BCAA** link to download BCAA now. For instructions on installing BCAA, refer to the *BCAA Service Requirements* document posted at MySymantec.

4. (Optional) If you have installed and configured a second BCAAA server for failover, enter the **Alternate server host** and **Port** values in the **Servers** section.

Note: If you plan to secure communication between the appliance and BCAAA, use a host name rather than an IP address. The DNS server that the appliance is configured to use must be able to resolve the hostname.


5. (Optional) In the **SSL Options** area, select **SSL enable** to enable SSL. Select the SSL device profile that this realm uses to make an SSL connection to the BCAAA server. You can choose any device profile that displays in the drop-down list. For information on using device profiles, see ["Appliance Certificates and SSL Device Profiles"](#) on page 1290.
6. Specify the type of credentials to accept from the browser/user agent. By default, all credential types are allowed and the ProxySG appliance will try to use Kerberos (the default authentication method for Windows clients), but will automatically downgrade to a different challenge type depending on the browser/user agent capabilities.
 - **Allow Basic credentials**—Prompts the user for a username and password to authenticate the user against the Windows Active Directory. Because the username and password are sent in plaintext, it is important to enable SSL between BCAAA and the ProxySG appliance if you allow Basic.

Note: Basic credentials cannot be disabled in the IWA realm if the IWA realm is part of a sequence realm but is not the first realm in the sequence with **try IWA authentication only once** enabled.

- **Allow NTLM credentials**—Uses an encrypted challenge/response that includes a hash of the password. Because the plaintext username and password are not sent across the wire, this method is more secure than Basic authentication.
- **Allow Kerberos credentials**—Uses a ticket containing an encrypted session key in place of a user name and password. This is the most secure method because it establishes mutual authentication between the client and the server using an encrypted shared key. However, if you select this option, NTLM is automatically selected as well; in the event that the browser/user agent and/or the BCAAA server are not configured properly for Kerberos, the appliance will automatically downgrade to NTLM. To use Kerberos, you must complete some additional configuration tasks. See ["Enabling Kerberos in a BCAAA Deployment"](#) on page 1017 for details.

Note: Forms authentication modes cannot be used with an IWA realm that allows only NTLM/Kerberos credentials. If a form mode is in use and the authentication realm is an IWA realm, you will receive a configuration error.

7. (Optional) To change the amount of time the appliance will wait for an authentication response from BCAAA before timing out, enter a new value in the **Timeout request after x seconds** field (default 60 seconds).
8. click **Apply**.
9. To verify that you have configured the realm successfully:
 - a. Click **Test Configuration**.
 - b. When prompted, enter the username and password of a user in the Windows domain and then click **OK**.
 - c. The appliance sends an authentication request to the configured server and then displays a message indicating whether the authentication succeeded or failed. If the test failed, go back and make sure you have configured the realm properly. If the test succeeds, the message also displays a list of any groups of interest (that is, groups that are referenced in policy) to which the user belongs.



Configuring a Direct Connection to the Windows Domain

If you have joined your ProxySG appliance to your Windows domain and created a realm for connecting to your Active Directory directly, you can configure and verify the authentication settings as follows:

1. Select the **Configuration > Authentication > IWA > IWA Servers** tab.


2. From the **Realm name** drop-down list, select the IWA realm you want to configure. If you have not yet created a realm, see ["Creating an IWA Realm"](#) on page 1018.
3. Specify the type of credentials to accept from the browser/user agent. By default, all credential types are allowed and the ProxySG appliance will try to use Kerberos (the default authentication method for Windows clients), but will automatically downgrade to a different challenge type depending on the browser/user agent capabilities.
 - **Allow Basic credentials**—Prompts the user for a username and password to authenticate the user against the Windows Active Directory.

Note: Basic credentials cannot be disabled in the IWA realm if the IWA realm is part of a sequence realm but is not the first realm in the sequence with **try IWA authentication only once** enabled.


- **Allow NTLM credentials**—Uses an encrypted challenge/response that includes a hash of the password.
- **Allow Kerberos credentials**—Uses a ticket containing an encrypted session key in place of a user name and password. This is the most secure method because it establishes mutual authentication between the client and the server using an encrypted shared key. However, if you select this option, NTLM is automatically selected as well; in the event that the browser/user agent and/or the ProxySG are not configured properly for Kerberos, the appliance will automatically downgrade to NTLM. To use Kerberos, you must complete some additional configuration tasks. See ["Enabling Kerberos in an IWA Direct Deployment"](#) on page 1016 for details.

Note: Forms authentication modes cannot be used with an IWA realm that allows only NTLM/Kerberos credentials. If a form mode is in use and the authentication realm is an IWA realm, you receive a configuration error.

4. (Optional) If you are sharing a service principal name (SPN) across multiple ProxySG appliances in a load balancing configuration, click **Set credentials**, enter the **User name** and **Password** for an Active Directory account, and then click **OK**. For details, see ["Using IWA Direct in an Explicit Kerberos Load Balancing/Failover Scenario"](#) on page 1034.



5. (Optional) To change the amount of time the appliance will wait for an authentication response before timing out, enter a new value in the **Timeout request after x seconds** field (default 60 seconds).
6. Click **Apply**.
7. To verify that you have configured the realm successfully:
 - a. Click **Test Configuration**.
 - b. When prompted, enter the username and password of a user in the Windows domain and then click **OK**.
 - c. The appliance sends an authentication request to the configured server and then displays a message indicating whether the authentication succeeded or failed. If the test failed, go back and make sure you have configured the realm properly. If the test succeeds, the message also displays a list of groups to which the user belongs.



Defining IWA Realm General Properties

Use the IWA General tab to configure the behavior of the authentication transaction, such as timeout and refresh intervals and cookie usage. You also use this tab to configure the Virtual URL for transparent authentication requests.

1. Select **Configuration > Authentication > IWA > IWA General**.

The screenshot shows the 'IWA General' configuration tab. It includes the following fields and options:

- Realm name:** A drop-down menu with 'IWA-Test' selected.
- Display name:** A text field containing 'IWA-Test'.
- Refresh Times:** A section with a checked checkbox 'Use the same refresh time for all'.
- Credential refresh time:** A text field with '900' and the unit 'seconds'.
- Surrogate refresh time:** A text field with '900' and the unit 'seconds'.
- Inactivity timeout:** A text field with '900' and the unit 'seconds'.
- Rejected credentials time:** A text field with '1' and the unit 'seconds'.
- Cookies:** A section with two checkboxes: 'Use persistent cookies' (unchecked) and 'Verify the IP address in the cookie' (checked).
- Virtual URL:** A text field containing 'www.cfauth.com/'.
- Challenge user after logout:** A checked checkbox.

2. From the **Realm name** drop-down list, select the IWA realm you want to configure. If you have not yet created a realm, see ["Creating an IWA Realm"](#) on page 1018.
3. (Optional) By default, the ProxySG appliance displays the authentication realm name when prompting the user for authentication credentials. To change the name that is displayed when the ProxySG appliance challenges the user for credentials from the default realm name, enter a new value in the **Display name** field, up to a maximum of 128 characters. This field cannot be left empty.
4. (Optional) If you want to change how often the appliance reauthenticates a client, modify the refresh and timeout values as follows:
 - **Credential refresh time**—(Basic credentials only) Specifies the amount of time the appliance will cache Basic credentials (username and password) and use these cached credentials to authenticate the user rather than sending another request to the authentication server. By default, basic credentials are good for 900 seconds (15 minutes).

- **Surrogate refresh time**—After the appliance successfully authenticates a client, it caches the client's IP address or a cookie (depending on the authentication mode that is in use) in its surrogate cache. If it receives subsequent requests from the same client during the surrogate refresh time, it uses the IP address or cookie in its cache to authenticate the user instead of sending a request to the authentication server. By default, the surrogate credential is good for 900 seconds (15 minutes).
- **Inactivity timeout**—When a client request is successfully authenticated, the appliance establishes an active session with the client and as long as that session stays active, the appliance will not attempt to reauthenticate requests from that client. This setting specifies how long the client session can be inactive before the appliance terminates the session; subsequent requests from that client will require authentication. By default, the client can be inactive for 900 seconds (15 minutes).

Note: If the **Challenge user after logout** option is selected, the appliance will automatically challenge the client for credentials when the session becomes inactive. If you are using a challenge method that prompts the user for credentials, you may want to deselect this option.

- **Rejected credentials time**—(Basic credentials only) Specifies whether to cache failed authentication attempts (bad password, expired account, disabled account, old password, or server down). If the client attempts to connect again during the rejected credentials time, the appliance will automatically reject the request for the specified period of time. Enter a value from 1 second (the default) to 10 seconds. Or, to disable this option, enter 0.
5. (optional) Modify how the appliance uses cookie surrogates by modifying the Cookies settings. These settings are only applicable if you plan to use an authentication mode that uses cookie surrogates.
- **Use persistent cookies**—By default, this option is deselected, which means that the appliance will use session cookies when creating a cookie surrogate for a client. Session cookies are only valid during the current browser session and are deleted when the user closes the browser. Therefore, the ProxySG appliance must reauthenticate the client each time the user starts a new browser session. If you select this option, the appliance will use persistent cookies instead of session cookies. Persistent cookies are stored on the client system and are therefore not deleted at the end of the browser session. When using persistent cookies, the appliance will only need to reauthenticate a client when the cookie in its surrogate credential database expires.
 - **Verify the IP address in the cookie**—By default, this option is selected, which means that the appliance will only accept a cookie from a client if the client IP matches the IP address in the surrogate cookie. To enable the appliance to accept cookies from IP addresses that do not match the address in the cookie—for example if you use DHCP—deselect this option.

6. (Transparent proxy only) Specify the URL to which to redirect client requests that require authentication in the **Virtual URL** field. For best results, the virtual URL you specify must:
 - Contain a simple hostname that does not contain any dots (for example, use `http://myproxy` rather than `http://myproxy.acme.com`. This allows IE to recognize the URL as part of the Intranet zone rather than the Internet zone so that the browser will automatically return credentials when challenged rather than prompting the user.
 - Resolve to the IP address of the ProxySG appliance. To accomplish this, you must add an "A" record to your internal DNS server that associates the Virtual URL with the IP address of the ProxySG appliance.
 - (IWA Direct Kerberos only) If you're using Kerberos in a non-load balancing IWA Direct realm, the Virtual URL must be the DNS name of the ProxySG appliance in the Active Directory domain. Typically this will be the DNS name of the Active Directory domain prefixed with the ProxySG appliance machine account name. For example, `sg.blue9.local`. If you do not use the Active Directory DNS name of the ProxySG as the Virtual URL, all authentication transactions will be downgraded to NTLM.
7. (Optional) If you want to prompt the client for authentication credentials whenever the inactivity timeout expires, select the **Challenge user after logout** check box.
8. Click **Apply**.

Creating the IWA Authentication and Authorization Policies

After you configure IWA on the ProxySG appliance (and set up BCAA, if applicable to your deployment), you must create the policy that instructs the appliance how to authenticate client requests. You can create a basic authentication policy that simply requires all requests to be authenticated and allows or denies access upon successful authentication. Or you can define more complex policies with rules that apply to a specific source address, subnet, port, user agent, or request header. You can even define different rules for different destinations. You can also create policies that allow access to guest users.

You can additionally create authorization policies that restrict access by user or group membership.

The following sections provides instructions for creating basic IWA authentication and authorization policies:

- ❑ ["Creating an IWA Authentication Policy"](#) on page 1028
- ❑ ["Creating a Guest Authentication Policy"](#) on page 1030
- ❑ ["Creating an IWA Authorization Policy"](#) on page 1031

Creating an IWA Authentication Policy

This section describes how to create a policy using the Visual Policy Manager (VPM). You can also create policy using the Content Policy Language (CPL).

Note that you must create an IWA realm before you can define the corresponding authentication policy.

1. Launch the VPM.
 - a. From the Management Console, select **Configuration > Policy > Visual Policy Manager**.
 - b. Click **Launch**.
2. Create the policy rule that enables the appliance to authenticate client requests:
 - a. Select **Policy > Add Web Authentication Layer**.
 - b. Enter a **Layer Name** or accept the default name and then click **OK**. The first policy rule displays with default settings.
3. Configure the authentication policy settings:
 - a. In the **Action** column of the first row, right-click and then select **Set**. The Set Action Object dialog displays.
 - b. Click **New** and then select one of the following authentication objects:
 - **Authenticate**—Use this option if you do not need to log user IDs for denied requests. With this option, if policy causes a request to be denied before the user is authenticated, the user ID associated with the request will not be available for access logging.
 - **Force Authenticate**—Use this option to ensure that user IDs are available for access logging (including denied requests).

Note: If you plan to create a guest authentication policy, create a combined object that contains the **Authenticate** object and a **Permit Authentication Error** object (be sure to select **All Except User Credentials Required**).

- c. (optional) Specify a **Name** for the authentication object.
- d. Select the **IWA Realm** from the drop-down list.
- e. Select the authentication mode from **Mode** drop-down list. Although you can select **Auto** to have the ProxySG appliance automatically choose an authentication mode, it is usually better to make a selection that is appropriate for your deployment as follows:
 - **Explicit deployments**—Select **Proxy** or **Proxy IP**. The Proxy IP mode reduces the load on the network because it uses an IP surrogate to reauthenticate clients that have already successfully authenticated.

- **Transparent deployments**—Select **Origin Cookie Redirect**. This mode redirects the client to the Virtual URL for authentication and uses a cookie surrogate to reauthenticate clients that have already successfully authenticated. The appliance will automatically downgrade to the Origin IP Redirect mode for user agents that do not support cookies.
- f. Click **OK** to close the Add Authenticate Object or Add Force Authenticate object dialog.
 - g. Click **OK** to close the Set Action Object dialog.
4. (optional) Restrict authentication to a subset of client requests, based on source or destination request attributes. The default settings in the policy rule will cause the ProxySG appliance to authenticate all client requests. You can set the Source and/or Destination columns to restrict authentication to a specified subset of requests. For example:
 - a. In the **Source** or **Destination** column of the first row, right-click and then select **Set**. The Set Source Object or Set Destination object dialog displays.
 - b. Click **New** and then select an object that represents the subset of requests you want to authenticate. After you select an object, you will be prompted to provide details. For example, if you choose the Client IP Address/Subnet object, you will be prompted for an IP address and subnet mask/prefix to which this rule will apply. When you first deploy your authentication policy, you may want to limit authentication to the source address of a test workstation or subnet. This allows you to identify and troubleshoot any configuration issues before rolling the policy out into production.
 5. (optional) Add additional policy rules to refine your authentication policy. A single Web Authentication Layer rule with the authenticate action is all you need to enable authentication. However, there may be some cases where you want to bypass authentication for certain requests and enable it for others. For example, you may have a certain client, subnet, or URL on which you do not require authentication or you may have some custom applications that do not know how to handle authentication requests. In this case, you would add an additional rule to your Web Authentication Layer policy to instruct the ProxySG appliance how to handle the exceptions. For example:
 - a. Click **Add Rule**. A new row appears in the Web Authentication Layer.
 - b. Specify which client requests this rule applies to by setting the **Source** or **Destination** columns.
 - c. Specify what the ProxySG appliance should do with requests that match the source and/or destination setting you have defined by right-clicking in the **Action** column of the row, selecting **Set**.
 - If you want to authenticate requests that match the specified source and/or destination request settings you have defined, click **New** and select **Authenticate** and click **OK**.

- If you want to bypass authentication for the matching requests, select **Do Not Authenticate** and click **OK**.
 - d. Arrange the rules according to how you want the ProxySG appliance to enforce them by selecting the rule you want to move and clicking **Move up** or **Move down**. The ProxySG appliance evaluates the rules in the order in which they appear in the policy layer. As soon as it finds a rule that matches the request, it will enforce the specified action (in this case, either to authenticate or not authenticate the request). Therefore, you should put more specific rules in front of general rules. For example, if you have a two rules in your policy—one that is set to authenticate requests from any source or destination and one that is set to not authenticate requests from a specific subnet—you would put the one that bypasses authentication in front of the general rule that matches all requests.
6. Install the authentication policy:
 - a. Click **Install policy**.
 - b. Click **OK** to acknowledge that the policy was successfully installed.

Creating a Guest Authentication Policy

A guest authentication policy enables users who do not have a Windows domain account on your network to access Internet resources.

Note: If you use guest authentication, remember that IWA/NTLM realms retrieve authorization data at the same time as the user is authenticated. In some cases, the system can distinguish between an authentication and authorization failure. Where the system cannot determine if the error was due to authentication or authorization, both the authentication and authorization are considered to be failed.

To create an IWA guest authentication policy:

1. Launch the VPM.
 - a. From the Management Console, select **Configuration > Policy > Visual Policy Manager**.
 - b. Click **Launch**.
2. Create a Web Authentication Layer for authenticating client requests for your domain users as described in "[Creating an IWA Authentication Policy](#)" on page 1028.
3. Create a second Web Authentication Layer to provide guest access:
 - a. Select **Policy > Add Web Authentication Layer**.
 - b. Enter a **Layer Name** to distinguish this layer from the previous layer (for example, Guest Authentication) and then click **OK**. The first policy rule displays with default settings.

4. Configure the source:
 - a. In the **Source** column of the first row, right-click and then select **Set**. The Set Source Object dialog displays.
 - b. Click **New** and then select **User Authentication Error**. The Add User Authentication Error Object dialog displays.
 - c. Select **Any errors** and click **OK** twice to save the source object and close the dialogs.
5. Configure the action:
 - a. In the **Action** column of the first row, right-click and then select **Set**. The Set Action Object dialog displays.
 - b. Click **New** and then select the **Authenticate Guest** object. The Add Authenticate Guest object dialog displays.
 - c. Select **Use realm** and then select your IWA realm from the drop-down list.
 - d. Enter a **Guest Username**. This will be the name that appears in your access log whenever guest access is granted; it does not correlate to an Active Directory user account.
 - e. Click **OK** twice to save the Action object and close the dialogs.
6. Make sure that the Web Authentication Layer for your guest policy is positioned after the your main Web Authentication Layer. To re-order the layers, select **Edit > Reorder Layers**.
7. Install the authentication policy:
 - a. Click **Install policy**.
 - b. Click **OK** to acknowledge that the policy was successfully installed.

Creating an IWA Authorization Policy

One of the benefits of IWA is that it automatically returns authorization information for a user in response to an authentication request. You do not have to perform any additional configuration to get authorization to work. After successfully authenticating a user, the appliance receives a list of all groups (IWA Direct) or groups of interest (IWA BCAA) to which the user belongs.

This section describes how to create a policy using the Visual Policy Manager (VPM). You can also create policy using the Content Policy Language (CPL).

1. Launch the VPM.
 - a. From the Management Console, select **Configuration > Policy > Visual Policy Manager**.
 - b. Click **Launch**.
2. Create a Web Access Layer:
 - a. Select **Policy > Add Web Access Layer**.
 - b. Enter a **Layer Name** or accept the default name and then click **OK**.

3. Specify the user or group to authorize (the source):
 - a. In the **Source** column of the first row, right-click and then select **Set**. The Set Source Object dialog displays.
 - b. Click **New** and then select the type of Active Directory object this rule will authorize:
 - To create a rule for authorizing a group, select **Group**. The Add Group Object dialog displays.
 - To create rule for authorizing a user, select **User**. The Add User Object dialog displays.
 - c. Select the IWA realm from the **Authentication Realm** drop-down list.
 - d. Specify the name of the Active Directory user or group that rule will authorize:
 - If you know the name of the Active Directory user or group, enter it in the **Group** or **User** field.
 - If you don't know the Active Directory name of the user or group, click **Browse** and select the group from the IWA Browser.
 - e. Click **OK** to close the Add Group Object or Add User Object dialog.
 - f. Click **OK** to close the Set Source Object dialog.
4. Specify whether to allow or deny requests from the specified user or group:
 - a. Right-click the **Action** column.
 - b. Select one of the following options:
 - **Allow**—Select this option if the default proxy policy for the appliance is set to deny proxy access through the ProxySG appliance. (This is the default in a secure web gateway deployment.)
 - **Deny**—Select this option of the default proxy policy for the appliance is set to allow proxy transactions. (This is the default in an acceleration deployment.)

If you aren't sure what the default proxy policy is set to on your appliance, go to **Configuration > Policy > Policy Options**.
5. (optional) Define any additional parameters that you want this rule to enforce.
6. To create additional authorization rules, repeat Steps 3 through 5.
7. Click **Install policy**.
8. Click **OK** to acknowledge that the policy was successfully installed.

Configuring Client Systems for Single Sign-On

One of the main benefits of IWA is that it can provide a single sign-on experience for users because it uses the workstation login to authenticate users. When configured properly, the browser will provide the credentials to the ProxySG appliance transparently when challenged for NTLM or Kerberos credentials (the user will always be prompted for Basic authentication credentials).

IWA only works with Windows domain credentials. If users log in to the workstation using local credentials instead of domain credentials, they will always be prompted whenever the ProxySG appliance returns an authentication challenge.

Both Internet Explorer (IE) and Firefox can be configured to provide authentication credentials to the ProxySG appliance transparently. By default, IE will automatically provide authentication credentials to any site in the local Intranet zone. If the Virtual URL for your ProxySG appliance contains a single hostname (that is, `http://myproxy` instead of `http://myproxy.acme.com`) you will not have to configure IE for IWA. If your Virtual URL does not fall within the Intranet zone, you will need to configure the IE to trust the URL. Firefox does not provide a single sign-on user experience for IWA by default and will therefore always need to be configured for single sign-on.

For explicit proxy deployments, you must also make sure the browser is configured to send requests to the ProxySG appliance. See ["About the Explicit Proxy"](#) on page 99 for details.

The procedure for configuring the browser to automatically provide login credentials to the ProxySG appliance is browser specific:

- ❑ ["Configure Internet Explorer for Single Sign-On"](#) on page 1033
- ❑ ["Configure Firefox for Single Sign-On"](#) on page 1034

Configure Internet Explorer for Single Sign-On

To configure IE for single-sign on with IWA:

1. Select **Tools > Internet Options**.
2. Select the **Security** tab.
3. Select the **Local intranet zone** and click **Sites > Advanced**.
4. Enter the fully qualified domain name of the ProxySG appliance (for explicit deployments) or the virtual URL (for transparent deployments) in the **Add this website to the zone** field and then click **Add > Close > OK**.
5. Select the **Advanced** tab and make sure the **Security > Enable Integrated Windows Authentication** option is selected.
6. Click **OK** to save your changes and close the Internet Options dialog.

Configure Firefox for Single Sign-On

To configure Firefox for single-sign on with IWA:

1. In the browser's **Location** field, enter `about:config`.
2. Click **I'll be careful, I promise!** to continue to the `about:config` page.
3. To get the browser to trust the ProxySG appliance and negotiate authentication with it, you must set values for the following options:
`network.automatic-ntlm-auth.trusted-uris`,
`network.negotiate-auth.delegation-uris`, `network.negotiate-auth.trusted-uris`. For each option, complete the following steps:
 - a. Locate the option you want to set by scrolling or entering the option name in the **Filter** field.
 - b. Double-click the option to open the Enter string value dialog.
 - c. Enter the fully qualified domain name of the ProxySG appliance (for explicit deployments) or the Virtual URL (for transparent deployments). If you have more than one ProxySG appliance that will challenge users for authentication credentials, separate the entries with commas.
4. Click **OK** to save your settings.

Using IWA Direct in an Explicit Kerberos Load Balancing/Failover Scenario

In a standard IWA Direct Kerberos deployment, the Kerberos service principal name (SPN) of the appliance is the appliance's own Active Directory machine account name. However, in a load balancing configuration, multiple ProxySGs must be able to decrypt the service tickets from the clients. For this reason, all ProxySGs in a load balancing group must share the same SPN. This will not work if each appliance uses its own machine account to process Kerberos authentication requests. In this case, you must create a new Active Directory account and use it to create a SPN that can be used by all appliances in the group. To deploy Kerberos in this configuration you must:

1. Set up a load balancing device in front of your appliances and designate a virtual IP address to use for all explicit proxy request. The load balancing device will then forward the requests to the ProxySGs in the group based on the load balancing rules you have defined.
2. Create a DNS entry for the device that resolves to this IP address. Note that the DNS name that you use must not map to an existing machine account name in Active Directory or the ProxySG appliance will not be able to authenticate Kerberos service tickets and authentication will fail.
3. Create an Active Directory account for the Kerberos load balancing user. This account does not need any special privileges. You will create the SPN using this account and the ProxySGs will use the account credentials to decrypt the service tickets from clients.
4. Use the Active Directory account you just created to create an SPN for the load balancing group as follows:

- a. Open a command prompt as administrator on the Domain Controller.
- b. Enter the following command:

```
setspn -A HTTP/<Load_Balancer_FQDN> <AD_Account_Name>
```


where *<Load_Balancer_FQDN>* is the fully qualified domain name (FQDN) of the load balancing device and *<AD_Account_Name>* is the name of the Active Directory user you created for the load balancing group. Note that this command is case-sensitive.

For example, if the FQDN of the load balancing device is `lb.acme.com` and the Active Directory account name you created is `KerberosLBUser`, you would enter the following command:

```
setspn -A HTTP/lb.acme.com KerberosLBUser
```

Note: Do not assign the same SPN to multiple Active Directory accounts or the browser will fall back to NTLM without providing any warning or explanation. To list all SPNs that are currently registered on an account, use the `setspn -L <AD Account Name>` command. If you find a duplicate, remove the extraneous SPN using the `setspn -D <SPN>` command.

5. On each ProxySG, create an IWA Direct realm. When configuring the realm on each appliance, you must provide the credentials for the AD Kerberos load balancing user you created. On the **IWA Servers** tab click **Set credentials**, enter the AD account **User name** and **Password**, and then click **OK**. Note that the user name you provide must be in the User Principal Name (UPN) format, for example `admin@acme.com`.



6. Configure the client browser explicit proxy settings to point to the FQDN of the load balancing device.

Chapter 54: Kerberos Constrained Delegation

This section discusses how to set up a realm to use Kerberos Constrained Delegation to provide authorized users with authenticated access to an OCS (Origin content server).

This section includes information about the following topics:

- ❑ ["About Kerberos Constrained Delegation"](#) on page 1037
- ❑ ["Symantec Implementation of Kerberos Constrained Delegation"](#) on page 1037
- ❑ ["KCD Process Overview"](#) on page 1038
- ❑ ["Requirements"](#) on page 1039
- ❑ ["Enabling Kerberos Constrained Delegation"](#) on page 1039
- ❑ ["Creating Kerberos Constrained Delegation Policies"](#) on page 1040
- ❑ ["Creating the CPL"](#) on page 1042

About Kerberos Constrained Delegation

Kerberos Constrained Delegation (KCD) offers a secure and reliable method of single sign on within Microsoft Windows networks. KCD is a Microsoft extension to the Kerberos protocol which enables a trusted process to acquire Kerberos tickets for a user without having access to that user's password. A single Kerberos Ticket authenticates a specific user to a specific service or server. KCD limits a process to only acquire tickets for users to a preconfigured set of services or servers.

Kerberos Constrained Delegation authentication cannot be used to authenticate administrative users to the ProxySG appliance management console.

Note: The ProxySG can handle extended Kerberos tickets, such as 32k authentication tokens. Note that the Kerberos token and other request headers must fit within the maximum size of the HTTP request header (128k).

Symantec Implementation of Kerberos Constrained Delegation

Blue Coat's implementation of KCD uses the ProxySG to authenticate the user. After authentication to the ProxySG, a Windows 2003 Server running BCAA (Blue Coat Authentication and Authorization Agent), provides Kerberos tickets to the ProxySG, allowing authorized users secure access to various backend services. You can authenticate a user's identity with any existing authorization realm.

The following diagram illustrates service request process for KCD:




Figure 54–1 Kerberos Constrained Delegation service request process.

KCD Process Overview

All deployments of KCD follows a similar high-level procedure. To enable Kerberos Constrained Delegation, you must:

- ❑ Create and configure an authentication realm allowing users to authenticate to the ProxySG.
- ❑ Create and configure an IWA realm to handle Kerberos authentication.
- ❑ Install and configure the BCAAA agent on a Windows Server. BCAAA provides tickets to the ProxySG on behalf of the authenticated user.
- ❑ Create policies enabling specific requests and connections.

Note: Kerberos Constrained Delegation does not have CLI commands. Refer to the CLI commands for the relevant authentication realms in the *Command Line Interface Reference*.

Requirements

Kerberos authentication requires two names to function: the user name (user principal name) and the service name (service principal name of the OCS). By default the ProxySG autogenerates SPNs in a similar manner to how Microsoft Internet Explorer autogenerates the SPN of the server it is attempting to authenticate to. You can override default behavior by setting the SPN using the VPM object: **Add Kerberos Constrained Delegation Object**.

Kerberos Constrained Delegation requires running the Windows domain at a Windows 2003 functional level. Although Windows Server 2000 supports the Kerberos protocol, it does not support constrained delegation and the protocol transition extensions, both of which are necessary.

Enabling Kerberos Constrained Delegation

All deployments of Kerberos Constrained Delegation require authentication to the ProxySG using an authentication realm (pre-existing or newly created). After authentication, the user is given access to the OCS using Kerberos. Because any authentication realm can be used, there are many deployment variants; however, the procedural differences are minimal. As a result, there is one basic procedure to enable KCD on the ProxySG.

To enable Kerberos Constrained Delegation:

Step	Task/Requirements	Management Console	Reference Information
1.	Create and configure an authentication realm. - KCD requires a full username (user principal name) to function. (Optional) Determine a user's authorization data.	<i>Configuration > Authentication > Authentication Realm</i>	For general information about realms, see " Controlling User Access with Identity-based Access Controls " on page 898. For information about a specific authentication realms, see the corresponding section.
2.	Create and configure an IWA realm to handle Kerberos. - IWA Realm must use SSL to connect to the BCAAA server. - IWA Realm must provide a certificate that BCAAA can verify.	<i>Configuration > Authentication > IWA</i> <i>Configuration > SSL > Device Profiles > Profiles</i>	" Creating an IWA Realm " on page 1018 " Appliance Certificates and SSL Device Profiles " on page 1290

Step	Task/Requirements	Management Console	Reference Information
3.	Configure BCAA to use Kerberos Constrained Delegation - Configure BCAA to run under the Local System account (default). - The BCAA server must be trusted to delegate to specified services using an authentication protocol. The SPNs for the services must be specified. - Select Require the ProxySG to provide a valid certificate in order to connect during BCAA installation. If using an existing installation, edit <code>bcaaa.ini</code> and set the value of <code>VerifySG</code> to 1.	BCAA You can download the Blue Coat Authentication and Authorization Agent at MySymantec. Complete instructions are also available online at: https://www.symantec.com/support-center/getting-started	Refer to the <i>BCAAA Service Requirements</i> document posted at MySymantec.
4.	Create policy to enable constrained delegation.	VPM: Under Web Authentication Layer>Action>Kerberos Constrained Delegation	<i>Visual Policy Manager Reference</i> "Creating Kerberos Constrained Delegation Policies" on page 1040

Creating Kerberos Constrained Delegation Policies

Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes.

The policy procedure below assumes no existing policy layers. A properly set up Visual Policy Manager has many existing layers and policies with a logical order. For existing deployments, it will be necessary to add new actions to existing layers to enable KCD. Make sure you have thoroughly read and are familiar with creating policies before continuing.

Note: Refer to the *Visual Policy Manager Reference* for complete details about the VPM.



There are two VPM objects that enable Kerberos Constrained Delegation: **Add Kerberos Constrained Delegation Object** and **Do not use Kerberos Constrained Delegation**. Both objects exist in the **Web Authentication Layer** as an **Action**. **Do not use Kerberos Constrained Delegation** is a fixed action and needs no configuration.

The following example policy enables Kerberos Constrained Delegation and shows configurable options.


To create KCD policies:

1. Select the **Configuration > Policy > Visual Policy Manager** tab.


- Click **Launch**. The VPM launches in a separate window.
- Select **Policy > Add Web Authentication Layer**. An **Add New Layer** dialog displays.



- Enter a name that is easily recognizable and click **OK**. A new policy tab and rule display in the VPM manager window.



- Select **Action** under the new rule. Right click **Any > Set**. The **Set Action Object** window displays.
- Select **New > Kerberos Constrained Delegation** to add a new Kerberos object.



7. The **Add Kerberos Constrained Delegation Object** window allows you to configure KCD implementation.
 - a. In the **Name** field, enter a name for the object or leave as is to accept the default.
 - b. From the **Authentication Type** drop-down list, select **origin** or **proxy**. If you are authenticating to an upstream origin server, select **origin**. If you are authenticating to a proxy server, select **proxy**.
 - c. In the **IWA Realm** field, enter a valid IWA realm to use for Kerberos authentication.
 - d. (Optional) Enter the **Service Principal Name** to use for the OCS. The default SPN for the service is set to `http/<hostname>`. If a non-standard port is used for a service, use `http/<hostname>:<port>`
8. Click **OK**.
9. Click **OK** to return to the VPM.
10. Click the **Install Policy** button when finished adding policies.

Creating the CPL

Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file layers.

- ❑ Authenticate to an upstream server with Kerberos constrained delegation.

```
<proxy>
url.host.exact="images.company.com" \
server.authenticate.constrained_delegation(origin, iwa_realm_1)
```

- ❑ Authenticate to an upstream proxy with Kerberos constrained delegation.

```
<proxy>
```

```
url.host.exact="proxy.company.com" \  
server.authenticate.constrained_delegation(proxy, iwa_realm_2)
```

- ❑ Set the service principal name to use when authenticating to an upstream server with Kerberos constrained delegation.

```
<proxy>  
url.host.exact="images.company.com" \  
server.authenticate.constrained_delegation(origin, iwa_realm_1) \  
server.authenticate.constrained_delegation.spn(http/  
images.company.com)
```


Chapter 55: LDAP Realm Authentication and Authorization

This section discusses Lightweight Directory Access Protocol (LDAP), the mechanism allowing query of LDAP compatible directory services.

Topics in this Section

This section includes information about the following topics:

- ❑ ["LDAP Overview"](#)
- ❑ ["Creating an LDAP Realm on the ProxySG" on page 1046](#)
- ❑ ["Configuring LDAP Properties on the ProxySG" on page 1048](#)
- ❑ ["Configuring LDAP Servers" on page 1048](#)
- ❑ ["Defining LDAP Base Distinguished Names" on page 1050](#)
- ❑ ["Defining LDAP Search & Group Properties" on page 1052](#)
- ❑ ["Customizing LDAP Objectclass Attribute Values" on page 1056](#)
- ❑ ["Defining LDAP General Realm Properties" on page 1057](#)
- ❑ ["Creating LDAP Authentication Policies Using the VPM" on page 1060](#)
- ❑ ["Creating LDAP Authentication Policies Using the CPL" on page 1062](#)
- ❑ ["LDAP Access Logging" on page 1063](#)
- ❑ ["LDAP Attribute Substitutions" on page 1063](#)

LDAP Overview

Lightweight Directory Access Protocol (LDAP) is a client protocol used to access information stored in an LDAP-compatible directory service. It is the vehicle by which LDAP-enabled applications speak to one another. As a shared protocol, LDAP integrates compatible applications in your network to a single authentication interface. Any additions or changes made to information in the directory are available to authorized users, directory-enabled applications, devices, and ProxySGs. This central control gives administrators simplified application management.

LDAP authentication realms can be used to authenticate administrative users (read only and read/write) to the management console. To ensure that credentials are not sent in clear text, configure the LDAP realm to use TLS to secure the communication with the LDAP server.

Note: To ensure that only TLS is used to communicate with the LDAP Server, check **Enable SSL** in the **LDAP Server** configuration page and edit the SSL device profile configured on the LDAP Server configuration page in the management console.

About the Symantec LDAP Solution

The ProxySG uses existing directory-based authentication by passing log in requests to the directory service. By keeping authentication centralized on your directory, a security administrator will always know who is accessing network resources and can easily define user/group-based policies to control access through the appliance.

Symantec supports both LDAP v2 and LDAP v3, but recommends LDAP v3 because it supports additional authentication mechanisms. In LDAP v3, servers can return referrals to other servers back to the client, allowing the client to follow those referrals if desired.

Supported Directory Services

LDAP group-based authentication for the ProxySG can be configured to support any LDAP-compliant directory including:

- ❑ Microsoft Active Directory Server
- ❑ Novell NDS/eDirectory Server
- ❑ Netscape/Sun iPlanet Directory Server

How to Implement LDAP Authentication

Configuring the SGOS for LDAP authentication involves the following steps:

1. Create an LDAP realm on the ProxySG.
2. Configure LDAP properties on the ProxySG.
 - a. Configure LDAP server settings
 - b. Define LDAP Base Distinguished Names
 - c. Define Authorization and Group information
 - d. Define objectclass attributes on an LDAP entry
 - e. Configure general LDAP realm settings
3. Create policies on the ProxySG.

Creating an LDAP Realm on the ProxySG

This section discusses the following topics:

- ❑ ["About LDAP Realms"](#)
- ❑ ["Creating an LDAP Realm"](#) on page 1047

About LDAP Realms

An LDAP authentication realm authenticates and authorizes users to access services using either explicit proxy or transparent proxy mode. These realms integrate third-party vendors, such as LDAP, Windows, and Novell, with the Blue Coat operating system.

Creating an LDAP Realm

Realm creation requires knowledge of LDAP server type, server host information, and attribute type. This section describes realm configuration options and how to set up and add additional realms. For more information, see "[LDAP Overview](#)" on page 1045.

To create an LDAP realm:

1. Select the **Configuration > Authentication > LDAP > LDAP Realms** tab.
2. Click **New**. The Add LDAP Realm dialog displays.

3. In the **Realm Name** field, enter a realm name. The name can be 32 characters long and composed of alphanumeric characters and underscores. The name *must* start with a letter.
4. Configure the realm options:
 - a. From the **Type of LDAP server** drop-down list, select the specific LDAP server.
 - b. Specify the host and port for the primary LDAP server. The host must be entered. The default port number is **389**.
 - c. (Optional) The ProxySG automatically retrieves the default **User attribute type** when the user specifies the LDAP server type.

You can manually specify the user attribute type for a particular LDAP server. The following list shows which attribute each directory server uses to form a username:

- Microsoft Active Directory Servers: `sAMAccountName=`
- Novell NDS/eDirectory Server/Other: `cn=`
- Netscape/iPlanet Directory Server: `uid=`

- d. Click **OK** to close the dialog.
5. Click **Apply**.

Configuring LDAP Properties on the ProxySG

After an LDAP authentication realm is created, you must set LDAP realm properties according to your directory type. This involves selecting server type, security method, LDAP version, LDAP search properties, group information, and general properties as described in the following topics:

- ❑ "Configuring LDAP Servers" on page 1048
- ❑ "Defining LDAP Base Distinguished Names" on page 1050
- ❑ "Defining LDAP Search & Group Properties" on page 1052
- ❑ "Customizing LDAP Objectclass Attribute Values" on page 1056
- ❑ "Defining LDAP General Realm Properties" on page 1057

Configuring LDAP Servers

After you create an LDAP realm, use the **LDAP Servers** page to change the current default settings.

To edit LDAP server properties:

Default values exist. You do not need to change these values if the default settings are acceptable.

1. Select the **Configuration > Authentication > LDAP > LDAP Servers** tab.

The screenshot shows the 'LDAP Servers' configuration page. On the left, there are numbered callouts with arrows pointing to specific fields:

- 2a** points to the 'Realm name' dropdown menu, which currently shows 'LDAP_test'.
- 2b** points to the 'Type of LDAP server' dropdown menu, which currently shows 'Microsoft Active Directory'.
- 2c** points to the 'LDAP Protocol Version' dropdown menu, which currently shows '3'.
- 3** points to the 'Primary server host' text input field, which contains '10.103.1.75'.
- 4** points to the 'SSL Options' section, which includes an 'Enable SSL' checkbox (unchecked) and an 'SSL device profile' dropdown menu (set to 'default').
- 5** points to the 'LDAP Specific Settings' section, which includes a 'Timeout request after' field set to '60' seconds.
- 6** points to the 'Case sensitive' checkbox, which is unchecked.


Other visible elements include the 'Servers' section with 'Alternate server host' and 'Port: 389' fields, and a 'Test Configuration' button at the bottom.

2. Configure realm information:
 - a. From the **Realm Name** drop-down list, select the LDAP realm for which you want to change server properties.
 - b. From the **Type of LDAP server** drop-down list, select the specific LDAP server.

- c. From the **LDAP Protocol Version** drop-down list, select **v2** for LDAP v2 support. LDAP v3 is the default.

If you use LDAP v3, you can select **Follow referrals** to allow the client to follow referrals to other servers. (This feature is not available with LDAP v2.) The default is **Disabled**.

3. Specify the host and port for the primary LDAP server. The host must be entered. The default port number is **389**. If you enable SSL, change the port to an SSL listening port, such as port **636**.
(Optional) Specify the host and port for the alternate LDAP server.
4. (Optional) Configure SSL options:
 - a. Under **SSL Options**, select **Enable SSL** to enable SSL. This option is valid *only* for LDAP v3.
 - b. Select the SSL device profile that this realm uses to make an SSL connection to a remote system. You can choose any device profile that displays in the drop-down list. For information on using device profiles, see "[Appliance Certificates and SSL Device Profiles](#)" on page 1290.
5. (Optional) Change the timeout request for the server from its default of **60** seconds.
6. If the LDAP server is configured to expect case-sensitive usernames and passwords, select **Case sensitive**.
7. Click **Apply**.
8. Verify the LDAP configuration as follows:
 - a. Click **Test Configuration**. The Test Configuration dialog displays.



- b. Enter the **Username** and **Password** of a client in your LDAP realm and then click **OK**. The ProxySG appliance will use configuration you supplied to send an authentication request to the LDAP server and return the results as follows:
 - If the LDAP server settings are configured properly, a dialog will display indicating that the test succeeded.

- If the test does not succeed, check that the settings on the **LDAP Servers** tab are configured properly and then test the configuration again.

Note: You can also look up LDAP users and groups from the CLI using the `lookup-user` and `lookup-group` commands. Refer to the *SGOS 6.5 Command Line Interface Reference* for details.

9. Repeat the above steps for additional LDAP realms, up to a total of 40.

Defining LDAP Base Distinguished Names

The ProxySG allows you to specify multiple Base Distinguished Names (DNs) to search per realm, along with the ability to specify a specific branch of a Base DN.

A *Base DN* identifies the entry that is starting point of the search. You must specify at least one non-null base-DN for LDAP authentication to succeed.

You must enter complete DN's. See the table below for some examples of distinguished name attributes.

Table 55–1 Distinguished Name Attributes


DN Attribute Syntax	Parameter Description
<code>c=country</code>	Country in which the user or group resides. Examples: <code>c=US</code> , <code>c=GB</code> .
<code>cn=common name</code>	Full name of person or object defined by the entry. Examples: <code>cn=David Smith</code> , <code>cn=Administrators</code> , <code>cn=4th floor printer</code>
<code>dc=domain component</code>	Component name of a domain. Examples: <code>cn=David Smith</code> , <code>ou=Sales</code> , <code>dc=MyDomain</code> , <code>dc=com</code>
<code>mail=e-mail address</code>	User or group e-mail address.
<code>givenName=given name</code>	User's first name.
<code>l=locality</code>	Locality in which the user or group resides. This can be the name of a city, country, township, or other geographic regions. Examples: <code>l=Seattle</code> , <code>l=Pacific Northwest</code> , <code>l=King County</code> .
<code>o=organization</code>	Organization to which the user or group is a member. Examples: <code>o=Blue Coat Inc</code> , <code>o=UW</code> .
<code>ou=organizational unit</code>	Unit within an organization. Examples: <code>ou=Sales</code> , <code>ou=IT</code> , <code>ou=Compliance</code> .
<code>st=state or province</code>	State or province in which the user or group resides. Examples: <code>st=Washington</code> , <code>st=Florida</code> .
<code>userPassword=password</code>	Password created by a user.

Table 55–1 Distinguished Name Attributes (Continued)

DN Attribute Syntax	Parameter Description
<code>streetAddress=street address</code>	Street number and address of user or group defined by the entry. Example: <code>streetAddress=4240 North Mary Avenue, Sunnyvale, California 94085</code> .
<code>sn=surname</code>	User's last name.
<code>telephoneNumber=telephone</code>	User or group telephone number.
<code>title=title</code>	User's job title.
<code>uid=user ID</code>	Name that uniquely identifies the person or object defined by the entry. Examples: <code>uid=ssmith</code> , <code>uid=kjones</code> .

To define searchable LDAP base DN's:


1. Select the **Configuration > Authentication > LDAP > LDAP DN** tab.




2. From the **Realm name** drop-down list, select the LDAP realm for which you want to change DN properties.
3. In the **User attribute type** field, the ProxySG has entered the default user attribute type for the type of LDAP server you specified when creating the realm.
 - Microsoft Active Directory Servers: `sAMAccountName=`
 - Novell NDS/eDirectory Server/Other: `cn=`
 - Netscape/iPlanet Directory Server: `uid=`

If you entered information correctly when creating the realm, you do not need to change the User attribute type in this step. If you do need to change or edit the entry, do so directly in the field.


- Enter as many Base DN's as required for the realm. Assume, for example, that Example Corp has offices in New York and Lisbon, each with its own Base DN. A simplified directory information tree is illustrated below.



To specify entries for the **Base DN's** field, click **New**, enter the Base DN, and click **OK**. Repeat for multiple Base DN's. To search all of Sample_Company, enter `o` values:



To search the manufacturing organizations, rather than starting at the top, enter `ou` and `o` values.



You can add, edit, and delete Base DN's for an ProxySG to search. The ProxySG searches multiple DN's in the order listed, starting at the top and working down. Select an individual DN and move it up or down in the list with the **Promote** and **Demote** buttons.

- Click **Apply**.

Defining LDAP Search & Group Properties

After creating an LDAP realm, providing at least the required fields of the LDAP server for that realm, and defining base DN's for the realm, you must define authorization properties for each LDAP realm you created.

Note: Authorization decisions are completely handled by policy. The groups that the appliance looks up and queries are derived from the groups specified in policy in `group=` conditions, `attribute=` conditions, `ldap.attribute=` conditions and `has_attribute` conditions. If you do not have any of those conditions, then Blue Coat does not look up any groups or attributes to make policy decisions based on authorization.

This section discusses the following types of LDAP searches:

- ❑ *Anonymous* searches, which allows a user to perform an LDAP search without entering a distinguished name.

To set up an anonymous search, see ["Enabling Anonymous LDAP Searches"](#).

- ❑ *Authenticated* searches, which require a search user DN to function properly.

To set up an authenticated search, see ["Enabling Authenticated LDAP Realm Searches"](#) on page 1054.


Enabling Anonymous LDAP Searches

The anonymous search feature allows a user to perform an LDAP search without entering a distinguished name. The LDAP directory attributes available for an anonymous client are typically a subset of those available when a valid user distinguished name and password have been used as search credentials.

For more information, see ["Defining LDAP Search & Group Properties"](#) on page 1052.

To allow anonymous LDAP realm searches:

1. Select the **Configuration > Authentication > LDAP > LDAP Search & Groups** tab.



2. From the **Realm name** drop-down list, select an LDAP realm for which you want to specify authorization information.
3. To permit users to anonymously bind to the LDAP service, select **Anonymous Search Allowed**. For example, with Netscape/iPlanet Directory Server, when anonymous access is allowed, no username or password is required by the LDAP client to retrieve information.

Note: Some directories require a valid user to be able to perform an LDAP search; they do not allow *anonymous bind*. (Active Directory is one such example.) For these directories, you must specify a valid fully-qualified

distinguished username and the password that permits directory access privileges. (For example, `cn=user1,cn=users,dc=bluecoat,dc=com` is a possible fully-qualified distinguished name.)

4. The **Dereference level** field has four values—**always**, **finding**, **never**, **searching**—that allow you to specify when to search for a specific object rather than search for the object's alias. The default is **Always**.
5. Click **Apply**.

Enabling Authenticated LDAP Realm Searches

Authenticated LDAP realm searches require a search user DN to function properly.

Note: For Microsoft Active Directory, you must use the full name and not the login name.

To enforce user authenticated LDAP realm searches:

1. Select the **Configuration > Authentication > LDAP > LDAP Search & Groups** tab.

The screenshot shows the 'LDAP Search & Groups' configuration window. On the left, a vertical list of numbers 2 through 6 has arrows pointing to specific fields in the window. Number 2 points to the 'Realm name' dropdown menu, which currently shows 'LDAP_test'. Number 3 points to the 'Search' section header. Number 4 points to the 'Anonymous search allowed' checkbox, which is currently unchecked. Number 5 points to the 'Search user DN' text input field. Number 6 points to the 'Dereference aliases' dropdown menu, which is currently set to 'always'. Other visible elements include a 'Change Password' button and a 'Change the search user password' link.

2. From the **Realm name** drop-down list, select an LDAP realm for which you want to specify authorization information.
3. To enforce user authentication before binding to the LDAP service, deselect **Anonymous Search Allowed**.
4. Enter a user distinguished name in the **Search User DN** field. This username can identify a single user or a user object that acts as a proxy for multiple users (a pool of administrators, for example). A search user distinguished name can be up to 512 characters long.
5. You can set or change the search user password by clicking **Change Password**. The password can be up to 64 alphanumeric characters long.

Note: You might want to create a separate user (such as Blue Coat, for example) instead of using an Administrator distinguished name and password.

6. The **Dereference level** field has four values—**always**, **finding**, **never**, **searching**—that allow you to specify when to search for a specific object rather than search for the object's alias. The default is **Always**.
7. Click **Apply**.

To define LDAP realm group information properties:

1. Select the **Configuration > Authentication > LDAP > LDAP Search & Groups** tab.

2. From the **Realm name** drop-down list, select an LDAP realm for which you want to specify authorization information.
3. Enter Membership type and Membership attribute: The ProxySG enters defaults for the following LDAP directories:
 - Microsoft Active Directory:
Membership type: `user`
Membership attribute type: `memberOf`
 - Netscape/Sun iPlanet:
Membership type: `group`
Membership attribute type: `uniqueMember`
 - Novell NDS eDirectory
Membership type: `group`
Membership attribute type: `member`
 - Other
Membership type: `user`
Membership attribute type: `member`
4. Username type to lookup: Select either **FQDN** or **Relative**. Only one can be selected at a time.
 - **Relative** can only be selected in the membership type is **Group**.
 - **FQDN** indicates that the lookup is done only on the user object. **FQDN** can be selected when the membership type is either **Group** or **User**.

5. Nested LDAP: If the LDAP server you use does not natively support group membership tests of nested groups, you can select the **Nested LDAP** checkbox.

Note: When a group of interest referenced within policy is part of a loop, User Authorization results in **Access Denied(policy_denied)**. For example, a loop forms if the group member **Testgroup** has the nested group member **Testgroup2**, which in turn has the aforementioned **Testgroup** as a nested member.

When loops are removed from an LDAP server, the **Nested Groups Support** option must be disabled and then re-enabled for the ProxySG to re-fetch the correct group structure.

6. Nested group attribute: For **other**, **ad** and **nds**, the default attribute is **member**. For **iPlanet**, the attribute is **uniqueMember**.
7. **Group constraint filter:** Enter a search limiting clause to reduce the number of groups returned for an LDAP search. This feature is generally used only when the user wishes to limit the scope of a comparison due to a very large number of groups. Constraints must be valid LDAP search filters and are AND'd to the search filter when performing a group search.

Example 1: If you enter `(cn=p*)` into the **Group constraint filter** field, only groups starting with the letter P are returned.

Example 2: If you enter `(cn=proxy)` into the **Group constraint filter** field, only the **proxy** group is returned.

Note: The **Group constraint filter** functions only for local comparisons. To enable local group comparisons, go to "[Defining LDAP General Realm Properties](#)" on page 1057.

8. Click **Apply**.

Customizing LDAP Objectclass Attribute Values


The *objectclass* attributes on an LDAP object define the type of object an entry is. For example, a user entry might have an *objectclass* attribute value of *person* while a group entry might have an *objectclass* attribute value of *group*.

The *objectclass* attribute values defined on a particular entry can differ among LDAP servers. The *objectclass* attribute values are attribute values only, they are not DNs of any kind.

Currently, the *objectclass* attribute values are used by Blue Coat during a VPM browse of an LDAP server. If an administrator wants to browse the groups in a particular realm, the ProxySG searches the LDAP server for objects that have *objectclass* attribute values matching those in the group list and in the container list. The list of *objectclass* attribute values in the container list is needed so that containers that contain groups can be fetched and expanded correctly.

To customize LDAP objectclass attribute values:

1. Select the **Configuration > Authentication > LDAP > LDAP Objectclasses** tab.



2. From the **Realm name** drop-down list, select the LDAP realm whose objectclasses you want to modify.
3. From the **Object type** drop-down list, select the type of object: **container**, **group**, or **user**.
4. To create or edit an object for the specified objectclass, click **New** or **Edit**. (The only difference is whether you are adding or editing an objectclass value.)
5. Enter or edit the objectclass, and click **OK**.
6. Click **Apply**.

Defining LDAP General Realm Properties

The **LDAP General** page allows you to specify the display name, the refresh times, an inactivity timeout value, cookies, virtual URL, and group comparison method.

To configure general LDAP settings:

1. Select the **Configuration > Authentication > LDAP > LDAP General** tab.

2. Configure realm information:

- From the **Realm name** drop-down list, select the LDAP realm for which you want to change properties.
- If needed, give the LDAP realm a display name. The default value for the display name is the realm name. The display name cannot be greater than 128 characters and it cannot be null.

3. Configure refresh option:

- Select the **Use the same refresh time for all** check box if you would like to use the same refresh time for all.
- Enter the number of seconds in the **Credential refresh time** field. The Credential Refresh Time is the amount of time basic credentials (username and password) are kept on the ProxySG. This feature allows the ProxySG to reduce the load on the authentication server and enables credential spoofing. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, the ProxySG will authenticate the user supplied credentials against the cached credentials. If the credentials received do not match the cached credentials, they are forwarded to the authentication server in case the user password changed. After the refresh time expires, the credentials are forwarded to the authentication server for verification.

- c. Enter the number of seconds in the field. The Surrogate Refresh Time allows you to set a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the ProxySG authenticates the transaction. After the refresh time expires, the ProxySG verifies the user's credentials. Depending upon the authentication mode and the user-agent, this may result in challenging the end user for credentials.

The main goal of this feature is to verify that the user-agent still has the appropriate credentials.

- d. Enter the number of seconds in the **Authorization refresh time** field. The Authorization Refresh Time allows you to manage how often the authorization data is verified with the authentication realm. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.
4. Enter the number of seconds in the **Inactivity timeout** field to specify the amount of time a session can be inactive before being logged out.
 5. If you use Basic credentials and want to cache failed authentication attempts (to reduce the load on the authentication service), enter the number of seconds in the **Rejected Credentials time** field. This setting, enabled by default and set to one second, allows failed authentication attempts to be automatically rejected for up to 10 seconds. Any Basic credentials that match a failed result before its cache time expires are rejected without consulting the back-end authentication service. The original failed authentication result is returned for the new request. All failed authentication attempts can be cached: Bad password, expired account, disabled account, old password, server down. To disable caching for failed authentication attempts, set the **Rejected Credentials time** field to 0.
 6. Configure the cookies option:
 - a. Select the **Use persistent cookies** check box to use persistent browser cookies instead of session browser cookies.
 - b. Select the **Verify the IP address in the cookie** check box if you would like the cookies surrogate credentials to only be accepted for the IP address that the cookie was authenticated. Disabling this will allow cookies to be accepted from other IP addresses.
 7. You can specify a virtual URL. For more information on the virtual URL, see ["About Origin-Style Redirection"](#) on page 910.
 8. Select the **Challenge user after logout** check box if the realm requires the users to enter their credentials after they have logged out.

9. Select the group comparison search method. There are two compare methods:
 - **Local**—The local method performs compare operations on the ProxySG after retrieving the appropriate entries. Because the compares are performed locally, this method typically reduces load on the LDAP server.
 - **Server**—The server method queries the LDAP server for each compare operation. If there are a large number of compares to perform, it can result in significant server load.

Note: There is a minute possibility that **local** compares can produce differing results from **server** compares. If you suspect erroneous compare results, set to **server**.

10. Click **Apply**.

Creating LDAP Authentication Policies

The following sections describe how to create LDAP authentication policies:

- ❑ ["Creating LDAP Authentication Policies Using the VPM" on page 1060](#)
- ❑ ["Creating LDAP Authentication Policies Using the CPL" on page 1062](#)

Creating LDAP Authentication Policies Using the VPM


This section describes how to create LDAP policy attributes. Keep in mind that this is just one part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes.

The following example lists the options available when creating an LDAP attribute policy using the VPM. The VPM allows you to perform LDAP string comparisons and existence checks. These LDAP attribute comparisons are performed locally on the ProxySG.


Note: Refer to the *Visual Policy Manager Reference* for details about VPM.

To launch the VPM:

1. Select the **Configuration > Policy > Visual Policy Manager** tab.
2. Click **Launch**. The VPM launches in a separate window.



3. Add a valid policy layer. The **LDAP Attribute Object** exists in the **Admin Access**, **SSL Access**, **Web Access**, and **Forwarding** layers as **Source** objects. For example, to add an **SSL Access** layer, select **Policy > Add SSL Access Layer**. An **Add New Layer** dialog box appears.
4. Enter a name that is easily understandable and click **OK**. A new policy tab and rule will display.
5. Select **source** for the new rule. Right click on **Any** and select **Set**. The **Set Source Object** window displays.



6. Select **New > LDAP Attribute** to create a new LDAP attribute object.

7. In the **Name** field, enter a name for the object or leave as is to accept the default.
8. From the **Authentication Realm** drop-down list, select a specific LDAP realm or **<ALL>**. The default setting for this field is **<ALL>**.
9. In the **Attribute Name** field, enter a valid attribute.
10. Select an attribute test method.
 - a. Select **Attribute Exists** to check if the attribute exists in the user's entry.
 - b. Select **Attribute value match** to check if an attribute matches the **Value** field. There are five attribute value match methods: **Exact Match**, **Contains**, **At Beginning**, **At end**, and **RegEx**.

Note: A list count check and numeric check are only available through CPL. For information about these checks, refer to the *Content Policy Language Guide*.

11. Click **OK**. You can add additional objects if necessary.
12. Click **OK** to return to the VPM.
13. Click the **Install Policy** button when finished adding policies.

Creating LDAP Authentication Policies Using the CPL

This section describes how to create LDAP policy attributes. Keep in mind that this is just one part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file layers.

Be aware that the default policy condition for these examples is `allow`. The default policy condition on new SGOS 5.x or later systems running the Proxy Edition is `deny`.

- ❑ Every LDAP-authenticated user is allowed access to the ProxySG.

```
<Proxy>
  authenticate (LDAPRealm)
```

- ❑ Group membership is the determining factor in granting access to the ProxySG.

```
<Proxy>
  authenticate (LDAPRealm)
<Proxy>
  group="cn=proxyusers, ou=groups, o=myco"
  deny
```

- ❑ A subnet definition determines the members of a group, in this case, members of the Human Resources department.

```
<Proxy>
  authenticate(LDAPRealm)
<Proxy>
  Define subnet HRSubnet
    192.168.0.0/16
    10.0.0.0/24
  End subnet HRSubnet
  [Rule] client_address=HRSubnet
    url.domain=monster.com
    url.domain=hotjobs.com
    deny
.
.
.
  [Rule]
    deny
```

LDAP Access Logging

The Blue Coat ProxySG uses the following ELFF field syntax for access logging.

```
x-ldap-attribute(<name>)
```

When the user is authorized the named attribute is fetched. When access log records are created, this field will be substituted with the value of the named attribute.

You enable Access Logging from the **Configuration > Access Logging > General** page. For information about customizing access logging, see [Chapter 29: "Creating Custom Access Log Formats"](#) on page 649.

LDAP Attribute Substitutions

LDAP attributes can be used as substitutions. The LDAP substitution uses the following syntax:

```
$(ldap.attribute.<name>)
```

Use of this LDAP substitution in any subber makes the attribute `<name>` interesting to all LDAP realms. When a user's entry is processed, objects of interest are obtained and associated with the user's login object. Whenever a substitution value is required, it is retrieved from the user's login object. If a list has more than one object, the value of the resulting substitution is in a comma separated list. If the attribute does not exist, the string is empty.

Note: Attribute names are case-sensitive so special care must be taken when using the LDAP substitution.

You can use the substitution to provide the value of an attribute in a header that is sent to an upstream server as well as within exception pages.

Notes

If you use guest authentication/authorization, note that:

- ❑ LDAP realms provide split authorization, and it is possible to be successfully authenticated but have authorization fail.

- ❑ If the LDAP realm `validate authorized user` command is disabled and the user does not exist in the authorization realm, authorization is considered a success and the user is assigned to the default group if there is one configured and it is of interest to policy.
- ❑ Returned attributes that are stored within the user's authentication data must not exceed 7680 bytes, or an authorization error occurs.

Chapter 56: Novell Single Sign-on Authentication and Authorization

This section discusses the Novell Single Sign-on (SSO) realm, which is an authentication mechanism that provides single sign-on authentication for users that authenticate against a Novell eDirectory server.

Topics in this Section

This section includes information about the following topics:

- ❑ ["About Novell SSO Realms"](#) on page 1065
- ❑ ["Creating a Novell SSO Realm"](#) on page 1067
- ❑ ["Novell SSO Agents"](#) on page 1068
- ❑ ["Adding LDAP Servers to Search and Monitor for Novell SSO"](#) on page 1069
- ❑ ["Querying the LDAP Novell SSO Search Realm"](#) on page 1070
- ❑ ["Configuring Authorization"](#) on page 1071
- ❑ ["Defining Novell SSO Realm General Properties"](#) on page 1072
- ❑ ["Modifying the sso.ini File for Novell SSO Realms"](#) on page 1074
- ❑ ["Creating the CPL"](#) on page 1075
- ❑ ["Notes"](#) on page 1076

About Novell SSO Realms

The mechanism uses the Novell eDirectory Network Address attribute to map the user's IP address to an LDAP FQDN. Because the mechanism is based on the user's IP address, it only works in environments where an IP address can be mapped to a unique user.

A Novell SSO realm consists of the following:

- ❑ BCAAA service information
- ❑ Novell eDirectory information
- ❑ Authorization realm information
- ❑ General realm information.

The Novell eDirectory information consists of a ProxySG LDAP realm that points to the master Novell eDirectory server that it is to be searched and monitored for user logins (see [Chapter 55: "LDAP Realm Authentication and Authorization"](#) on page 1045 for information on configuring LDAP realms) and a list of eDirectory server and port combinations that specify additional servers

to monitor for logins. Additional monitor servers must be specified if they contain user information that is not replicated to the master Novell eDirectory server being searched.

After a Novell SSO realm has been configured, you can write policy that authenticates and authorizes users against the Novell SSO realm.

To ensure that users who do not successfully authenticate against the Novell SSO realm are not challenged, administrators can use a realm sequence that contains the Novell SSO realm and then a policy substitution realm to use when Novell SSO authentication fails.

Note: The Novell SSO realm works reliably only in environments where one IP address maps to one user. If an IP address cannot be mapped to a single user, authentication fails. Those with NAT systems, which uses one set of IP addresses for intranet traffic and a different set for Internet traffic, may need to use a different realm for authentication.

When a user logs into the Novell network, the user entry in Novell eDirectory is updated with the login time and the IP address that the user logged in from and the login time. The ProxySG uses BCAAA to do LDAP searches and monitoring of the configured Novell eDirectory servers to obtain the user login information and maintain a user IP address to user FQDN map.

To create the initial IP/FQDN map, the BCAAA service searches the configured master eDirectory server for all user objects within the configured base DN that have a Network Address attribute. For each user entry returned, BCAAA parses the Network Address attribute and adds the IP/FQDN entry to the map. If an existing entry exists for that IP address, it is overwritten.

A user entry can have more than one Network Address entry in which case an entry for each IP address is added to the map. Since service accounts can login using the same IP address and subsequently overwrite entries for actual users, the BCAAA service has a configurable list of the Service names to ignore. Users can be added or removed from the list in the sso.ini file. (see ["Modifying the sso.ini File for Novell SSO Realms"](#) on page 1074.)

Once the initial map has been created it is kept current by monitoring all of the eDirectory servers that contain unique partition data for the eDirectory tree. By default, the search server defined by the LDAP realm is monitored. If other servers contain data that is not replicated to the search server, they must be individually monitored. When a server is being monitored, each time a user logs in or logs out, an event message is sent to BCAAA to update its mapping of FQDNs to IP addresses.

Multiple ProxySG devices can talk to the same BCAAA service and can reference the same eDirectory servers. To avoid multiple queries to the same server, the LDAP hostname and port combination uniquely identifies an eDirectory configuration and should be shared across devices.

To ensure that BCAAA has complete map of FQDNs to IP addresses, the realm can be configured to do a full search of the configured master eDirectory server up to once per day.

The BCAA service must be version 120 or higher and must be installed on a machine that can access the eDirectory server. The BCAA machine does not need to have a Windows trust relationship with the eDirectory server.

Note: Refer to the *BCAAA Service Requirements* document for up-to-date information on BCAA compatibility. The *BCAAA Service Requirements* document is posted at MySymantec.

A Novell SSO realm can be configured to perform no authorization, authorize against itself (the default), or authorize against another valid authorization realm. When a Novell SSO realm is configured to authorize against itself, authorization is done through the LDAP search realm specified by the Novell SSO realm. The behavior is similar to the Novell SSO realm explicitly selecting the LDAP realm as the authorization realm.


Novell SSO realms are compatible with administrative authentication configurations, but not recommended because they do not challenge the user to authenticate. Novell SSO relies on the LDAP server to identify the user requesting access based on their client IP address.

Creating a Novell SSO Realm

The **Configuration > Authentication > Novell SSO > Novell SSO Realms** tab allows you to create a new Novell SSO realm. Up to 40 Novell SSO realms can be created.

To Create a Novell SSO Realm through the Management Console

1. Select the **Configuration > Authentication > Novell SSO > Novell SSO Realms** tab.
2. Click **New**. The Add Novell SSO Realm dialog displays.



3. In the **Realm name** field, enter a realm name. The name can be 32 characters long and composed of alphanumeric characters and underscores. The name *must* start with a letter.
4. Click **OK** to close the dialog.
5. Click **Apply**.

Novell SSO Agents

You must configure the Novell realm so that it can find the Blue Coat Authentication and Authorization Agent (BCAAA).

Novell SSO Agent Prerequisite

You must have defined at least one Novell SSO realm (using the Novell SSO Realms tab) before attempting to configure the BCAAA agent. If the message **Realms must be added in the Novell SSO Realms tab before editing this tab** is displayed in red at the bottom of this page, you do not currently have any Novell SSO realms defined.

1. Select the **Configuration > Authentication > Novell SSO > Agents** tab.

The screenshot shows the 'Novell SSO Agents' configuration window. The 'Agents' tab is active. The 'Realm name' is set to 'NSSO_test'. The 'Primary agent' section shows a host of '192.168.20.1' and a port of '16101', with buttons to 'Change Private Key Password' and 'Change Public Certificate Password'. The 'Alternate agent' section has empty fields for host and port, also with password change buttons. The 'SSL Options' section has 'Enable SSL' checked and 'SSL device profile' set to 'default'. The 'Test Configuration' section includes a message about requiring an IP address and a 'Test Configuration' button. The 'Novell SSO Specific Settings' section shows a 'Timeout request after' of '60' seconds.

2. Select the realm name to edit from the drop-down list.
3. In the **Primary Agent** section, enter the hostname or IP address where the BCAAA agent resides. Change the port from the default of **16101** if necessary.
(Optional) You can change the encrypted passwords for the private key and public certificate on the BCAAA machine that are to be used for SSL communication between the BCAAA service and the Novell eDirectory server by clicking **Change Private Key Password** or **Change Public Certificate Password**. The location of the private key and public certificate are specified in the `sso.ini` file on the BCAAA machine. (For information on changing the location of the private key and public certificate, see ["Modifying the sso.ini File for Novell SSO Realms"](#) on page 1074.)
4. (Optional) Enter an alternate agent host and agent name in the **Alternate agent** section. As with the Primary Agent, you can change the passwords for the private key and public certificate for the alternate agent.

The primary and alternate BCAAA server must work together to support fail-over. If the primary BCAAA server fails, the alternate server should be able to search and monitor the same set of eDirectory servers.

5. (Optional) Configure SSL options:
 - a. Click **Enable SSL** to enable SSL between the ProxySG and the BCAAA.
 - b. (Optional) Select the SSL device profile that this realm uses to make an SSL connection to a remote system. You can choose any device profile that displays in the drop-down list. For information on using device profiles, see "[Appliance Certificates and SSL Device Profiles](#)" on page 1290.

Note: The **Enable SSL** setting only enables SSL between the ProxySG and BCAAA. To enable SSL between BCAAA and the eDirectory server, the **Enable SSL** setting must be set in the LDAP search realm.

6. In the **Timeout Request** field, enter the number of seconds the ProxySG allows for each request attempt before timing out. (The default request timeout is **60** seconds.)
7. Click **Apply**.
8. Verify the Novell SSO configuration as follows:
 - a. Click **Test Configuration**. The Test Configuration dialog displays.
 - b. Enter the **IP address** of a client system in your Novell Directory and then click **OK**. The ProxySG appliance will use configuration you supplied to send an authentication request to BCAAA and return the results as follows:
 - If the ProxySG and the BCAAA server are configured properly, BCAAA will return the LDAP DN of the user associated with the IP address you provided.
 - If the test does not succeed, check that the settings on the **Agents** tab as well as the BCAAA settings are configured properly and then test the configuration again.

Adding LDAP Servers to Search and Monitor for Novell SSO

The BCAAA service searches and monitors specified eDirectory servers to determine which users are logged in and their Network Address attribute value. Those attribute values are converted into IP addresses, and BCAAA maintains a map of IP addresses to LDAP FQDNs.


If the eDirectory tree is partitioned across multiple servers, the realm must monitor every eDirectory server that has unique user information.

LDAP Server Prerequisite

You must have defined at least one Novell SSO realm (using the Novell SSO Realms tab) before attempting to specify LDAP server configuration. If the message **Realms must be added in the Novell SSO Realms tab before editing this tab** is displayed in red at the bottom of this page, you do not currently have any Novell SSO realms defined.

To specify the eDirectory servers:

1. Select the **Configuration > Authentication > Novell SSO > LDAP Servers** tab.



2. Select the realm name to edit from the drop-down list.
3. Select an LDAP realm from the drop-down list. The servers configured in this LDAP realm are used to do the full searches of the eDirectory tree.
4. If you have a deployment with multiple servers holding partitions that are not fully replicated to the master server, you can monitor each LDAP server individually.
 - a. To add an LDAP server to monitor, click **New**.
 - b. Add the IP address and port of the LDAP server and click **OK** to close the dialog.
 - c. Repeat for additional LDAP servers you need to monitor.
5. Click **Apply**.

Querying the LDAP Novell SSO Search Realm

You can specify the time and days that a full search of the eDirectory tree is repeated in order to ensure that the mappings maintained by BCAA are up to date.

LDAP Novell SSO Search Real Prerequisite

You must have defined at least one Novell SSO realm (using the Novell SSO Realms tab) before attempting to configure LDAP queries. If the message **Realms must be added in the Novell SSO Realms tab before editing this tab** is displayed in red at the bottom of this page, you do not currently have any Novell SSO realms defined.

To specify search criteria:

1. Select the **Configuration > Authentication > Novell SSO > LDAP Queries** tab.

2. Select the realm name to edit from the drop-down list.
3. In the full search pane, specify the time of day you want the search to take place from the drop-down list.
4. Select or clear check boxes to specify days to search.
5. If you have changed the **Novell eDirectory Network Address** or **Login Time LDAP** attribute name, you can enter those changed names in the **Network Address LDAP name** and the **Login Time LDAP name** fields. The names must match the LDAP names configured on the eDirectory server for authentication to succeed.
6. Click **Apply**.

Configuring Authorization

Novell SSO realm can be configured to do no authorization, authorize against itself (the default), or authorize against another valid authorization realm (either LDAP or Local).

Authorization Prerequisite

You must have defined at least one Novell SSO realm (using the Novell SSO Realms tab) before attempting to configure authorization. If the message **Realms must be added in the Novell SSO Realms tab before editing this tab** is displayed in red at the bottom of this page, you do not currently have any Novell SSO realms defined.

To specify an authorization realm:

1. Select the **Configuration > Authentication > Novell SSO > Authorization** tab.

2. From the Realm Name drop-down list, select the Novell SSO realm to edit.
3. By default, the Novell SSO realm is selected to authorize against itself by default. To select another realm, clear the **Self** check box and select an authorization realm from the drop-down list.
4. The LDAP FQDN is selected as the **Authorization user name**, by default. Change this if the user's authorization information resides in a different root DN. To select a different authorization name, clear the **Use FQDN** option and enter a different name. For example:
`cn=$(user.name),ou=partition,o=company`
5. Click **Apply**.

Defining Novell SSO Realm General Properties

The **Novell SSO General** tab allows you to specify the refresh times, an inactivity timeout value, and cookies, and a virtual URL.

Novell SSO realms default to the **origin-ip** authentication mode when no authentication mode or the auto authentication mode is specified in policy. After a user has first successfully authenticated to the ProxySG, all subsequent requests from that same IP address for the length of the surrogate credential refresh time are authenticated as that user. If the first user is allowed or denied access, subsequent users during that same time coming from the same IP address are allowed or denied as that first user. This is true even if policy would have treated them differently if they were authenticated as themselves.

If multiple users often log in from the same IP address, it is recommended to use a shorter surrogate credential refresh timeout than the default or an authentication mode that does not use IP surrogate credentials.

Novell SSO Prerequisite

You must have defined at least one Novell SSO realm (using the Novell SSO Realms tab) before attempting to set Novell SSO general properties. If the message **Realms must be added in the Novell SSO Realms tab before editing this tab** is displayed in red at the bottom of this page, you do not currently have any Novell SSO realms defined.

To configure Novell SSO general settings:

1. Select the **Configuration > Authentication > Novell SSO > Novell SSO General** tab.

2. From the **Realm name** drop-down list, select the Novell SSO realm for which you want to change properties.
3. Configure refresh options:
 - a. Select the **Use the same refresh time for all** option to use the same refresh time for all.
 - b. Enter the number of seconds in the **Surrogate refresh time** field. The Surrogate Refresh Time allows you to set a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the ProxySG authenticates the transaction. After the refresh time expires, the ProxySG determines which user is using the current IP address, and update the surrogate credential to authenticate with that user.
 - c. Enter the number of seconds in the **Authorization refresh time** field. The Authorization Refresh Time allows you to manage how often the authorization data is verified with the authentication realm. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.
4. Type the number of seconds in the **Inactivity timeout** field to specify the amount of time a session can be inactive before being logged out.
5. Configure cookie options:
 - a. Select the **Use persistent cookies** check box to use persistent browser cookies instead of session browser cookies.

- b. Select the **Verify the IP address in the cookie** check box if you would like the cookies surrogate credentials to only be accepted for the IP address that the cookie was authenticated. Disabling this allows cookies to be accepted from other IP addresses.
6. You can specify a virtual URL. For more information on the virtual URL, see ["About Origin-Style Redirection"](#) on page 910.
7. Click **Apply**.

Modifying the sso.ini File for Novell SSO Realms

The Novell SSO realm uses the `sso.ini` file for configuration parameters required by the BCAAA service to manage communication with the Novell eDirectory server. Three sections in the `sso.ini` file are related to the Novell SSO realm: `NovellSetup`, `NovellTrustedRoot Certificates`, and `SSOServiceUsers`. You only need to modify settings in the `NovellTrustedRoot Certificates` section if the LDAP realm used by the Novell SSO realm requires that the identity of the server be verified.

The `sso.ini` file is located in the BCAAA installation directory.

Note: The changes to the `sso.ini` file have no effect until the BCAAA service is restarted.

To modify Novell SSO realms parameters:

1. Open the file in a text editor.
2. In the Novell Setup section, modify the parameters as needed (the default values are as follows):
 - `MonitorRetryTime=30`
 - `SearchRetryTime=30`
 - `TrustedRootCertificateEncoding=der`
 - `PublicCertificateEncoding=der`
 - `PrivateKeyFile=`
 - `PrivateKeyEncoding=der`
3. If the LDAP realm used by the Novell SSO realm requires that the identity of the server be verified, add the paths to the `Trusted root certificate` files in the `NovellTrustedRootCertificates` section.
4. In the `SSOServiceUsers` section, list the names of users who can log in with eDirectory credentials on behalf of the service and mask the identity of the logged-on user.

Listing these users here forces the BCAAA service to ignore them for authentication purposes.
5. Save the `sso.ini` file.

Creating the CPL

You can create CPL policies now that you have completed Novell SSO realm configuration. Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes.

Note: The examples below assume the default policy condition is `allow`.

Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file layers.

- ❑ Every Novell SSO-authenticated user is allowed access the ProxySG.

```
<Proxy>
  authenticate (NSSOR realm)
```

- ❑ Group membership is the determining factor in granting access to the ProxySG.

```
<Proxy>
  authenticate (NSSOR realm)
<Proxy>
  group="cn=proxyusers, ou=groups, o=myco" ALLOW
deny
```

Using Single Sign-On Realms and Proxy Chains

Some Application Delivery Network (ADN) configurations mask the source IP address of the request. For example, if the path for a request is:

client workstation > branch proxy > data center proxy > gateway proxy

policy running on the gateway might see the IP address of the data center proxy rather than the IP address of the client workstation.

Note: The source IP address is not masked if you use the **reflect client ip** attribute.

In this ADN configuration, policy must be configured so that Windows SSO, Novell SSO, and policy substitution realms authenticate users correctly.

Use the `user.login.address` and `authenticate.credentials.address` policy gestures to override the IP address of the credentials used for authentication and match the IP address of the authenticated user.

Note: The `user.login.address` condition only works correctly if you use the `authenticate.credentials.address` property to set the address.

You can also use the `x-cs-user-login-address` substitution to log this event.

Examples

In the following example, the address to use for authenticating with **myrealm** is set to the address received from the `HTTP Client-IP` header.

```
<proxy>
  authenticate(myrealm) \
  authenticate.credentials.address($(request.header.Client-IP))
```

In the following example, the user is authenticated if logged in from the 1.2.3.0/24 subnet.

```
<proxy>
  user.login.address=1.2.3.0/24 allow
```

Notes

- ❑ The Novell SSO realm works reliably only in environments where one IP address maps to one user. NAT environments are not supported.
- ❑ Novell SSO realms are not supported in IPX environments.
- ❑ Event monitoring of eDirectory is only compatible with eDirectory 8.7+.
- ❑ Upgrade to Novell client 4.91 SP1 or later if you experience issues with the Network Address attribute not being updated during login.
- ❑ Novell SSO realms do not use user credentials so they cannot spoof authentication information to an upstream server.
- ❑ If an upstream proxy is doing Novell SSO authentication, all downstream proxies must send the client IP address.
- ❑ There can be response time issues between the BCAAA service and the eDirectory servers during searches; configure the timeout for LDAP searches to allow the eDirectory server adequate time to reply.

Chapter 57: Policy Substitution Realm

This section describes Policy Substitution realms, which provide a mechanism for identifying and authorizing users based on information in the request to the ProxySG. It includes the following topics:

- ❑ ["About Policy Substitution Realms"](#)
- ❑ ["Creating a Policy Substitution Realm" on page 1080](#)
- ❑ ["Configuring User Information" on page 1081](#)
- ❑ ["Creating a List of Users to Ignore" on page 1083](#)
- ❑ ["Configuring Authorization" on page 1083](#)
- ❑ ["Defining Policy Substitution Realm General Properties" on page 1084](#)
- ❑ ["Creating the Policy Substitution Policy" on page 1086](#)

About Policy Substitution Realms

The Policy Substitution realm is used typically for best-effort user discovery, mainly for logging and subsequent reporting purposes, without the need to authenticate the user. Be aware that if you use Policy Substitution realms to provide granular policy on a user, it might not be very secure because the information used to identify the user can be forged.

The realm uses information in the request and about the client to identify the user. The realm is configured to construct user identity information by using policy substitutions.

Substitution Realms are not compatible with administrative authentication to the ProxySG appliance management console.

If authorization data (such as group membership) is required, configure the realm with the name of an associated authorization realm (such as LDAP or local). If an authorization realm is configured, the fully-qualified username is sent to the authorization realm's authority to collect authorization data.

You can use policy substitutions realms in many situations. For example, a Policy Substitution realm can be configured to identify the user:

- ❑ based on the results of a NetBIOS over TCP/IP query to the client computer.
- ❑ based on the results of a reverse DNS lookup of the client computer's IP address.
- ❑ based on the contents of a header in the request. This might be used when a downstream device is authenticating the user.
- ❑ based on the results of an Ident query to the client computer.

The realm is configured the same way as other realms, except that the realm uses policy substitutions to construct the username and full username from information available in and about the request. Any policy substitution whose value is available at client logon can be used to provide information for the name.

The Policy Substitution realm, in addition to allowing you to create and manipulate realm properties (such as the name of the realm and the number of seconds that credential cache entries from this realm are valid) also contains attributes to determine the user's identity. The user's identity can be determined by explicitly defining the usernames or by searching a LDAP server. The following two fields are used to determine the user's identity by definition:

- ❑ A user field: A string containing policy substitutions that describes how to construct the simple username.
- ❑ A full username field: A string containing policy substitutions that describes how to construct the full username, which is used for authorization realm lookups. This can either be an LDAP FQDN when the authorization realm is an LDAP realm, or a simple name when local realms are being used for authorization.

Note: The user field and username field must include at least one substitution that successfully evaluates in order for the user to be considered authenticated.

If no policy substitutions exist that map directly to the user's simple and full usernames but there are substitutions that map to attributes on the user on the LDAP server, the user's identity can be determined by searching the LDAP server. The following fields are used to determine the user's identity by LDAP search:

- ❑ LDAP search realm: The LDAP realm on the ProxySG that corresponds to the LDAP server where the user resides
- ❑ Search filter: An LDAP search filter as defined in RFC 2254 to be used in the LDAP search operation. Similar to the explicitly defined username and full username fields, the search filter string can contain policy substitutions that are available based on the user's request. The search filter string must be escaped according to RFC 2254. The policy substitution modifier `escape_ldap_filter` is recommended to use with any policy substitutions that could contain characters that need to be escaped. It will escape the policy substitution value per RFC 2254.

Note: The search filter must include at least one substitution that successfully evaluates before the LDAP search will be issued and the user authenticated.

- ❑ **User attribute:** The attribute on the search result entry that corresponds to the user's full username. If the search result entry is a user entry, the attribute is usually the FQDN of that entry. The user's full username is the value of the specified attribute. If the attribute value is an FQDN, the user's simple username is the value of the first attribute in the FQDN. If the attribute value is not an FQDN, the simple username is the same as the full username.

Note: Policy Substitution realms never challenge for credentials. If the username and full username cannot be determined from the configured substitutions, authentication in the Policy Substitution realm fails.

Remember that Policy Substitution realms do not require an authorization realm. If no authorization realm is configured, the user is not a member of any group. The effect this has on the user depends on the authorization policy. If the policy does not make any decisions based on groups, you do not need to specify an authorization realm. Also, if your policy is such that it works as desired when all Policy Substitution realm users are not in any group, you do not have to specify an authorization realm.

After the Policy Substitution realm is configured, you must create policy to authenticate the user.

Note: If all the policy substitutions fail, authentication fails. If any policy substitution works, authentication succeeds in the realm.

Example

The following is an example of how to use substitutions with Policy Substitution realms.

Assumptions:

- ❑ The user `susie.smith` is logged in to a Windows client computer at IP address `10.25.36.47`.
- ❑ The Windows messenger service is enabled on the client computer.
- ❑ The client computer is in the domain `AUTHTEAM`.
- ❑ The customer has an LDAP directory in which group information is stored. The DN for a user's group information is
`cn=username,cn=users,dc=computer_domain,dc=company,dc=com`
where `username` is the name of the user, and `computer_domain` is the domain to which the user's computer belongs.
- ❑ A login script that runs on the client computer updates a DNS server so that a reverse DNS lookup for `10.25.36.47` results in
`susie.smith.authteam.location.company.com`.

Results:

Under these circumstances, the following username and full username attributes might be used:

- ❑ **Username:** `$(netbios.messenger-username)@$(client.address).`

This results in `SUSIE.SMITH@10.25.36.47.`

- ❑ **Full username:** `cn=$(netbios.messenger-username),cn=users,dc=$(netbios.computer-domain),dc=company,dc=com.`

This results in `cn=SUSIE.SMITH,cn=users, dc=AUTHTEAM,dc=company,dc=com.`

- ❑ **Username:** `$(netbios.computer-domain)\$(netbios.messenger-username).`

This results in `AUTHTEAM\SUSIE.SMITH.`

- ❑ **Username:** `$(client.host:label(6)).$(client.host:label(5)).`

This results in `SUSIE.SMITH.`

Example

The following is an example of how to determine the user's identity by search.

Assumptions:

- ❑ The user `susie.smith` is logged in to a Windows client computer.
- ❑ The customer has an LDAP directory in which group information is stored. The FQDN for Susie Smith is `cn=Susie Smith, cn=Users, dc=Eng, dc=company, dc=com.`

Results:

Under these circumstances the login username can not be explicitly mapped to the user's FQDN, so a search of the LDAP server for the user's login identity is required instead. The following values can be used:


- ❑ **Search filter:** `(sAMAccountName=$(netbios.messenger-username:escape_ldap_filter))`
- ❑ **User attribute:** default of FQDN

This results in a simple username of `Susie Smith` and a full username of `cn=Susie Smith, cn=Users, dc=Eng, dc=company, dc=com.`

Creating a Policy Substitution Realm

To create a Policy Substitution realm:

1. Select the **Configuration > Authentication > Policy Substitution > Policy Substitution Realms** tab.
2. Click **New**; the Add Policy Substitution Realm dialog displays.



3. In the **Realm name** field, enter a realm name. The name can be up to 32 characters long and composed of alphanumeric characters and underscores. The name *must* start with a letter.
4. Click **OK** to close the dialog.
5. Click **Apply**.

Configuring User Information

This section describes how to add user search information.

Prerequisites

You must have defined at least one Policy Substitution realm (using the **Policy Substitution Realms** tab) before attempting to set Policy Substitution realm properties. If the message `Realms must be added in the Policy Substitutions Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any Policy Substitution realms defined.

To Define Policy Substitution User Information:

1. Select the **Configuration > Authentication > Policy Substitution > User Information** tab.

2 → Realm name: PolicySub_test

3 → ☒ Determine username by definition

-or-

4 → ☐ Determine username by search

Username: \$(netbios.messenger-username)@\$(client.address).

Full username: ,cn=users, dc=\$(netbios.computer-domain),dc=company,dc=com

LDAP search realm name: <None>

Search filter:

User attribute: ☐ FQDN

2. From the **Realm name** drop-down list, select the Policy Substitution realm for which you want to change realm properties.
3. To determine username by definition, select **Determine username by definition** and specify the username and full username strings. Remember that the **Username** and **Full username** attributes are character strings that contain policy substitutions. When authentication is required for the transaction, these character strings are processed by the policy substitution mechanism, using the current transaction as input. The resulting string becomes the user's identity for the current transaction. For an overview of usernames and full usernames, see ["About Policy Substitution Realms"](#) on page 1077.
- or-
4. To determine username by search, select **Determine username by search**.
 - From the drop-down list, select the LDAP realm to use as a search realm.
 - The search filter must be a valid LDAP search filter per RFC 2254. The search filter can contain any of the policy substitutions that are available based on the user's request (such as IP address, netbios query result, and ident query result).
 - The user attribute is the attribute on the LDAP search result that corresponds to the user's full username. The LDAP search usually results in user entries being returned, in which case the user attribute is the FQDN. If the LDAP search was for a non-user object, however, the username might be a different attribute on the search result entry.
5. Click **Apply**.

Creating a List of Users to Ignore

This section describes how to create a list of users to be ignored during an LDAP username search (see ["Configuring User Information"](#) on page 1081).

Prerequisite

You must have defined at least one Policy Substitution realm (using the Policy Substitution Realms tab) before attempting to set Policy Substitution realm properties. If the message `Realms must be added in the Policy Substitutions Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any Policy Substitution realms defined.

1. Select **Configuration > Authentication > Policy Substitution > Ignore Users**.
2. From the **Realm Name** drop-down list, select the Policy Substitution realm for which you want to change realm properties.
3. Click **New** to add a username to be ignored during the username search. The username format depends on what the LDAP search is looking for but will most often be an LDAP FQDN.
4. Click **OK** to close the dialog; repeat the previous step to add other users.
5. Click **Apply**.

Configuring Authorization


Policy Substitution realms do not require an authorization realm. If the policy does not make any decisions based on groups, you need not specify an authorization realm.

Prerequisite

You must have defined at least one Policy Substitution realm (using the Policy Substitution Realms tab) before attempting to set Policy Substitution realm properties. If the message `Realms must be added in the Policy Substitutions Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any Policy Substitution realms defined.

To configure an authorization realm:

1. Select the **Configuration > Authentication > Policy Substitution > Authorization** tab.



2. From the **Realm Name** drop-down list, select the Policy Substitution realm for which you want to change realm properties.

3. From the **Authorization Realm Name** drop-down list, select the authorization realm you want to use to authorize users.
4. Click **Apply**.

Defining Policy Substitution Realm General Properties

The Policy Substitution General tab allows you to specify the refresh times, an inactivity timeout value, cookies, and a virtual URL.

Prerequisite

You must have defined at least one Policy Substitution realm (using the **Policy Substitution Realms** tab) before attempting to set Policy Substitution general properties. If the message `Realms must be added in the Policy Substitution Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any Policy Substitution realms defined.

To configure Policy Substitution realm general settings

1. Select the **Configuration > Authentication > Policy Substitution > General** tab.

The screenshot shows the 'General' tab of the Policy Substitution configuration interface. It contains the following fields and options:

- 2** → **Realm name:** A drop-down menu showing 'PolicySub_test'.
- 3** → **Refresh Times:** A section with a checked checkbox 'Use the same refresh time for all'. Below it are two fields: 'Surrogate refresh time: 900 seconds' and 'Authorization refresh time: 900 seconds'.
- 4** → **Inactivity timeout:** A field set to '900 seconds'.
- 5** → **Cookies:** A section with an unchecked checkbox 'Use persistent cookies' and a checked checkbox 'Verify the IP address in the cookie'.
- 6** → **Virtual URL:** A text field containing 'www.cfauth.com/'.

2. From the **Realm name** drop-down list, select the Policy Substitution realm for which to change properties.
3. Configure refresh options:
 - a. Select the **Use the same refresh time for all** check box if you would like to use the same refresh time for all.

- b. Enter the number of seconds in the **Surrogate refresh time** field. The Surrogate Refresh Time allows you to set a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the ProxySG authenticates the transaction. After the refresh time expires, the ProxySG reevaluates the user's credentials.
 - c. Enter the number of seconds in the **Authorization refresh time** field. The Authorization Refresh Time allows you to manage how often the authorization data is verified with the authentication realm. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.
4. Enter the number of seconds in the **Inactivity timeout** field to specify the amount of time a session can be inactive before being logged out.
5. Configure cookie options:
 - a. Select the **Use persistent cookies** option to use persistent browser cookies instead of session browser cookies.
 - b. Select the **Verify the IP address in the cookie** option if you would like the cookies surrogate credentials to only be accepted for the IP address that the cookie was authenticated. Disabling this will allow cookies to be accepted from other IP addresses.
6. You can specify a virtual URL. For more information on the virtual URL, see ["About Origin-Style Redirection"](#) on page 910.
7. Click **Apply**.

Notes

- ❑ Following are examples of how to configure four different types of Policy Substitution realms. For a list of available substitutions, see the *Content Policy Language Guide*.
- Identity to be determined by sending a NetBIOS over TCP/IP query to the client computer, and using LDAP authorization

```
SGOS#(config) security policy-substitution create-realm netbios
SGOS#(config) security policy-substitution edit-realm netbios
SGOS#(config policy-substitution netbios) username \
$(netbios.messenger-username)
SGOS#(config policy-substitution netbios) full-username \
cn=$(netbios.messenger-username),cn=users,dc=company,dc=com
SGOS#(config policy-substitution netbios) authorization-realm-name
ldap
```

- Identity to be determined by reverse DNS, using local authorization. Blue Coat assumes login scripts on the client computer update the DNS record for the client.

```
SGOS#(config) security policy-substitution create-realm RDNS
SGOS#(config) security policy-substitution edit-realm RDNS
SGOS#(config policy-substitution RDNS) username \
$(client.host:label(5)).$(client.host:label(6))
#SGOS#(config policy-substitution RDNS) full-username \
$(client.host:label(5)).$(client.host:label(6))
SGOS#(config policy-substitution RDNS) authorization-realm-name
local
```

- Identity to be determined by a header in the request, using LDAP authorization.

```
SGOS#(config) security policy-substitution create-realm header
SGOS#(config) security policy-substitution edit-realm header
SGOS#(config policy-substitution header) username \
$(request.x_header.username)
SGOS#(config policy-substitution header) full-username \
cn=$(request.x_header.username),cn=users,dc=company,dc=com
SGOS#(config policy-substitution header) username \ authorization-
realm-name ldap
```

- Identity to be determined by sending an Ident query to the client computer

```
SGOS#(config) security policy-substitution create-realm ident
SGOS#(config) security policy-substitution edit-realm ident
SGOS#(config policy-substitution ident) username $(ident.username)
SGOS#(config policy-substitution ident) full-username
"cn=$(ident.username),cn=Users,dc=company,dc=com"
```

- ❑ If you need to change the NetBIOS defaults of 5 seconds and 3 retries, use the `nbstat requester` option from the `netbios` command submode. (For more information on using the NetBIOS commands, refer to the *Command Line Interface Reference*.)
- ❑ If you need to change the Ident defaults of 30 second timeout, treating username whitespace as significant and querying Ident port 113, use the client commands in the `identd` command submode. (For more information on using the Ident commands, refer to the *Command Line Interface Reference*.)

Creating the Policy Substitution Policy

When you complete Policy Substitution realm configuration, you must create CPL policies for the policy-substitution realm to be used. Be aware that the example below is just part of a comprehensive authentication policy. By themselves, they are not adequate.

For policy substitution realms, the username and group values are case-sensitive.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file `<Proxy>` and other layers.

Be aware that the default policy condition for this example is `allow`. On new SGOS 5.x or later systems running the Proxy Edition, the default policy condition is `deny`. Every Policy Substitution realm authenticated user is allowed to access the ProxySG.

```
<Proxy>
  authenticate (PolicySubstitutionRealm)
```

Using Single Sign-On Realms and Proxy Chains

Some Application Delivery Network (ADN) configurations mask the source IP address of the request. For example, if the path for a request is:

client workstation > branch proxy > data center proxy > gateway proxy

policy running on the gateway might see the IP address of the data center proxy rather than the IP address of the client workstation.

Note: The source IP address is not masked if you use the **reflect client ip** attribute.

In this ADN configuration, policy needs to be configured so that Windows SSO, Novell SSO, and policy substitution realms can authenticate users correctly.

Use the `user.login.address` and `authenticate.credentials.address` policy gestures to override the IP address of the credentials used for authentication and match the IP address of the authenticated user.

Note: The `user.login.address` condition only works correctly if you use the `authenticate.credentials.address` property to set the address.

You can also use the `x-cs-user-login-address` substitution to log this event.

Examples

In the following example, the address to use for authenticating with **myrealm** is set to the address received from the HTTP Client-IP header.

```
<proxy>
  authenticate (myrealm) \
  authenticate.credentials.address ($ (request.header.Client-IP) )
```

In the following example, the user is authenticated if logged in from the 1.2.3.0/24 subnet.

```
<proxy>
  user.login.address=1.2.3.0/24 allow
```


Chapter 58: RADIUS Realm Authentication and Authorization

This section discusses RADIUS authentication and authorization.

Topics in this Section

This section includes information about the following topics:

- ❑ ["Creating a RADIUS Realm"](#) on page 1090
- ❑ ["Defining RADIUS Realm Properties"](#) on page 1090
- ❑ ["Defining RADIUS Realm General Properties"](#) on page 1092
- ❑ ["Creating the Policy"](#) on page 1094
- ❑ ["Troubleshooting"](#) on page 1097

About RADIUS

RADIUS is often the protocol of choice for ISPs or enterprises with very large numbers of users. RADIUS is designed to handle these large numbers through centralized user administration that eases the repetitive tasks of adding and deleting users and their authentication information. RADIUS also inherently provides some protection against sniffing.

Some RADIUS servers support one-time passwords. One-time passwords are passwords that become invalid as soon as they are used. The passwords are often generated by a token or program, although pre-printed lists are also used. Using one-time passwords ensures that the password cannot be used in a replay attack.

The ProxySG's one-time password support works with products such as Secure Computing SafeWord synchronous and asynchronous tokens and RSA SecurID tokens.

The ProxySG supports RADIUS servers that use challenge/response as part of the authentication process. SafeWord asynchronous tokens use challenge/response to provide authentication. SecurID tokens use challenge/response to initialize or change PINs.

The challenge is displayed as the realm information in the authentication dialog; Blue Coat recommends that you use form authentication if you create a challenge/response realm, particularly if you use SecurID tokens.

If you set an authentication mode that uses forms, the system detects what type of question is being asked. If it is a yes/no question, it displays the query form with a *yes* and *no* button. If it is a new PIN question, the system displays a form with entry fields for the new PIN.

For information on using form authentication, see [Chapter 65: "Forms-Based Authentication"](#) on page 1199.

Using policy, you can fine-tune RADIUS realms based on RADIUS attributes. If you use the Blue Coat attribute, groups are supported within a RADIUS realm.

RADIUS authentication is compatible with administrative authentication for the ProxySG appliance management console. Because the communication between the ProxySG appliance and the RADIUS server cannot be encrypted, this realm type is not recommended for administrator authentication.

Creating a RADIUS Realm

To create a RADIUS realm:

You can create up to 40 RADIUS realms.

1. Select the **Configuration > Authentication > RADIUS > RADIUS Realms** tab.
2. Click **New**. The Add RADIUS Realm dialog displays.

3. In the **Realm name** field, enter a realm name.
The name can be 32 characters long and composed of alphanumeric characters and underscores. The name *must* start with a letter.
4. Specify the host and port for the primary RADIUS server. The default port is **1812**.
5. Specify the RADIUS secret. RADIUS secrets can be up to 64 characters long and are always case sensitive.
6. Click **OK**.
7. Click **Apply**.

Defining RADIUS Realm Properties

Once you have created the RADIUS realm, you can change the primary host, port, and secret of the RADIUS server for that realm.

To re-define RADIUS server properties:

1. Select the **Configuration > Authentication > RADIUS > RADIUS Servers** tab.

The screenshot shows the 'RADIUS Servers' configuration page. At the top, there are tabs for 'RADIUS Realms', 'RADIUS Servers' (selected), and 'RADIUS General'. Below the tabs, the 'Realm name' is set to 'RADIUSTest'. The 'Primary Server' section has a 'Host' field with '192.158.9.9' and a 'Port' field with '1812', along with a 'Change Secret' button. The 'Alternate Server' section has empty 'Host' and 'Port' fields and a 'Change Secret' button. The 'Test Configuration' section contains a message: 'Validating the realm configuration requires the username and password of a user within the realm.' and a 'Test Configuration' button. The 'RADIUS Specific Settings' section includes a 'Send credentials to server encoded with character set' dropdown set to 'Unicode: UTF8', 'Timeout request after' and 'retry' fields both set to '5', and checkboxes for 'One-time passwords' (unchecked) and 'Case sensitive' (checked).

2. From the **Realm Name** drop-down list, select a RADIUS realm.

3. Specify the host and port for the primary RADIUS server.

The default port is **1812**. (To create or change the RADIUS secret, click **Change Secret**. RADIUS secrets can be up to 64 characters long and are always case sensitive.)

4. (Optional) Specify the host and port for the alternate RADIUS server.
5. From the **Send credentials to server encoded with character set** drop-down list, select the character set used for encoding credentials; the RADIUS server needs the same character set.

A character set is a Multipurpose Internet Mail Extension (MIME) charset name. Any of the standard charset names for encodings commonly supported by Web browsers can be used. The default is **Unicode:UTF8**.

6. In the **Timeout Request** field, enter the total number of seconds the ProxySG will attempt to connect to RADIUS servers; the contact to the other server occurs when half of the timeout period has lapsed. The default request timeout is 10 seconds.

In the **Retry** field, enter the number of attempts you want to permit before marking a server offline.

The client maintains an average response time from the server; the retry interval is initially twice the average. If that retry packet fails, then the next packet waits twice as long again. This increases until it reaches the timeout value. The default number of retries is **10**.

7. If you are using one-time passwords, select the **One-time passwords** option.

You must enable one-time passwords if you created a challenge/response realm.

8. If the RADIUS server is configured to expect case-sensitive usernames and passwords, make sure the **Case sensitive** option is selected.
9. Click **Apply**.
10. Verify the RADIUS configuration as follows:
 - a. Click **Test Configuration**. The Test Configuration dialog displays.
 - b. Enter the **Username** and **Password** of a client in your RADIUS realm and then click **OK**. The ProxySG appliance will use configuration you supplied to send an authentication request to the RADIUS server and return the results as follows:
 - If the RADIUS server settings are configured properly, a dialog will display indicating that the test succeeded. It will also display a list of groups to which the user belongs.
 - If the test does not succeed, check that the settings on the **RADIUS Servers** tab are configured properly and then test the configuration again.

Defining RADIUS Realm General Properties

The **RADIUS General** tab allows you to specify the display name, the refresh times, an inactivity timeout value, cookies, and a virtual URL.

To configure general settings:

1. Select the **Configuration > Authentication > RADIUS > RADIUS General** tab.

The screenshot shows the 'RADIUS General' configuration tab. It contains several fields and checkboxes. Numbered annotations (2-8) point to specific elements:

- 2** points to the 'Realm name' dropdown menu, which is set to 'RADIUSTest'.
- 3** points to the 'Display name' text field, which contains 'RADIUSTest'.
- 4** points to the 'Refresh Times' section, specifically the 'Use the same refresh time for all' checkbox, which is checked.
- 5** points to the 'Credential refresh time' field, which is set to '900' seconds.
- 6** points to the 'Surrogate refresh time' field, which is set to '900' seconds.
- 7** points to the 'Inactivity timeout' field, which is set to '900' seconds.
- 8** points to the 'Rejected credentials time' field, which is set to '1' seconds.

Below these fields is a 'Cookies' section with two checkboxes: 'Use persistent cookies' (unchecked) and 'Verify the IP address in the cookie' (checked). At the bottom, there is a 'Virtual URL' field set to 'www.cfauth.com/' and a 'Challenge user after logout' checkbox, which is checked.

2. Configure name options:

- a. From the **Realm name** drop-down list, select the RADIUS realm for which you want to change properties.
- b. (Optional) In the **Display Name** field, change the RADIUS realm display name.

The default value for the display name is the realm name. The display name cannot be greater than 128 characters and it cannot be empty.

3. Configure refresh options:

- a. Select the **Use the same refresh time for all** check box if you would like to use the same refresh time for all.
- b. Enter the number of seconds in the **Credential refresh time** field.

The Credential Refresh Time is the amount of time basic credentials (username and password) are kept on the ProxySG. This feature allows the ProxySG to reduce the load on the authentication server and enables credential spoofing. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, the ProxySG authenticates the user supplied credentials against the cached credentials. If the credentials received do not match the cached credentials, they are forwarded to the authentication server in case the user password changed. After the refresh time expires, the credentials are forwarded to the authentication server for verification.

- c. Enter the number of seconds in the **Surrogate refresh time** field.

The Surrogate Refresh Time allows you to set a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the ProxySG authenticates the transaction. After the refresh time expires, the ProxySG verifies the user's credentials. Depending upon the authentication mode and the user-agent, this may result in challenging the end user for credentials.

The main goal of this feature is to verify that the user-agent still has the appropriate credentials.

4. Type the number of seconds in the **Inactivity timeout** field to specify the amount of time a session can be inactive before being logged out.
5. If you use Basic credentials and want to cache failed authentication attempts (to reduce the load on the authentication service), enter the number of seconds in the **Rejected Credentials time** field.

This setting, enabled by default and set to one second, allows failed authentication attempts to be automatically rejected for up to 10 seconds. Any Basic credentials that match a failed result before its cache time expires are rejected without consulting the back-end authentication service. The original failed authentication result is returned for the new request.

All failed authentication attempts can be cached: Bad password, expired account, disabled account, old password, server down.

To disable caching for failed authentication attempts, set the **Rejected Credentials time** field to 0.

6. Configure cookie options:
 - a. Select the **Use persistent cookies** check box to use persistent browser cookies instead of session browser cookies.
 - b. Select the **Verify the IP address in the cookie** check box if you would like the cookies surrogate credentials to only be accepted for the IP address that the cookie was authenticated.

Disabling this allows cookies to be accepted from other IP addresses.

7. You can specify a virtual URL. For more information on the virtual URL, see ["About Origin-Style Redirection"](#) on page 910.
8. Select the **Challenge user after logout** check box if the realm requires the users to enter their credentials after they have logged out.
9. Click **Apply**.

Creating the Policy

Fine-tune RADIUS realms through attributes configured by policy—CPL or VPM. You can also create RADIUS groups. To configure RADIUS realm attributes, continue onto the next sections. To create RADIUS groups, see ["Creating RADIUS Groups"](#) on page 1096.

Note: RADIUS groups can only be configured through policy. This feature is not available through either the Management Console or the CLI.

Configuring RADIUS Realm Attributes

RADIUS Realm attributes can be configured using the `attribute.name` and `has_attribute.name` CPL conditions and source objects in VPM. For more information about policy and supported attributes, refer to these conditions in the *Content Policy Language Guide*.

Creating User-Defined RADIUS Attributes

You can also create user-defined RADIUS attributes using the CLI. If you plan on using the ProxySG as a session monitor and want the attributes available for use in a session monitor, you must reference the attributes to the session monitor as well. For more information about configuring the session monitor and referencing the attributes, see ["Configuring the ProxySG as a Session Monitor"](#) on page 1099.

Use the following CLI commands to create user-defined RADIUS attributes:

Table 58–1 User-defined RADIUS attribute command descriptions

Command	Options	Description
add radius-attribute	<radius-type (1-255)> <attribute name> [integer tag-integer ipv4 ipv6] [string tag-string] <max-length (1-247)> [<[enum tag-enum] (1-253)>=<string <max-length (1-253)>> { <(1-253)>=<string <max-length (1-253)>>}]	Add a new RADIUS attribute.
add vendor-attribute	<vendor id> <vendor-type (1-255)> <attribute name> [integer tag-integer ipv4 ipv6] [string tag-string] <max-length (1-247)> [<[enum tag-enum] (1-253)>=<string <max-length (1-253)>> { <(1-253)>=<string <max-length (1-253)>>}]	Add a vendor specific attribute.
remove	<attribute name>	Remove a RADIUS attribute. This will not remove attributes that are <i>currently</i> part of the session-monitor's configuration.

Examples: Configuring User-Defined RADIUS Attributes

The following examples describe how to configure user-defined RADIUS attributes.

Example 1

The following example shows an enum mapping an integer value to a string value:

```
SGOS#(config radius attributes) add radius-attribute 205 sample-enum
enum 1="string for value 1" 2=string2 3="string for value 3"
```

The integer values are sent on the wire from the RADIUS server. However, an admin can also refer to a value using either an integer or a string in CPL using the following expressions:

```
session-monitor.attribute.sample-enum=3
session-monitor.attribute.sample-enum="string for value 3"
```

Example 2

The following example shows octet string value:

```
SGOS#(config radius attributes) add radius-attribute 206 sample-octet-
string octet-string 30
```

An octet string functions similarly to a string, but it can contain binary data.

Example 3

The following example show a tag data type:

```
SGOS#(config radius attributes) add radius-attribute 205 sample-tag-string tag-string 25
```

Tag data types differ from non-tag counterparts because they include an extra byte in the value sent from the RADIUS server, which identifies a VPN tunnel. The ProxySG skips this extra value to get to the *actual* value when parsing the value sent from the RADIUS server.

Example 4

The following example shows a vendor attribute with a fictional vendor ID value of 21234:

```
SGOS#(config radius attributes) add radius-attribute 21234 1 sample-vendor-integer integer
```

Example 5

To safely modify the configuration of an existing RADIUS attribute, you must remove it from the system and add it again with the new configuration. The following example shows how to change the maximum length of the User-Name attribute.

1. Back up the ProxySG policy and install a new blank policy.
2. (If the attribute is in use in the session monitor) Remove the attribute from the RADIUS session monitor:

```
#(config session-monitor attributes) remove user-name
```

3. Remove the attribute from RADIUS configuration:

```
#(config radius attributes) remove user-name
```

4. (If you removed the attribute from the session monitor) Restart the appliance.
5. Add the User-Name attribute and specify the new length of 64 characters:

```
#(config radius attributes) add radius-attribute 1 user-name string 64
```

6. (If the attribute was previously in use in the session monitor) Add the attribute to the RADIUS session monitor:

```
#(config session-monitor attributes) add user-name
```

7. Restore the policy you backed up in step 1.

Creating RADIUS Groups

Create a RADIUS realm group by using the custom Blue Coat attribute, which can appear multiple times within a RADIUS response. It can be used to assign a user to one or more groups. Values that are found in this attribute can be used for comparison with the group condition in CPL and the group object in VPM. The group name is a string with a length from 1-247 characters. The Blue Coat Vendor ID is 14501, and the Blue-Coat-Group attribute has a Vendor Type of 1.

If you are already using the `Filter-ID` attribute for classifying users, you can use that attribute instead of the custom `Blue-Coat-Group` attribute. While the `Filter-ID` attribute does not work with the CPL group condition or the group object in VPM, the `attribute.Filter-ID` condition can be used to manage users in a similar manner.

CPL Example

The examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file layers.

- ❑ Every RADIUS-authenticated user is allowed access the ProxySG if the RADIUS attribute `service-type` is set.

```
<Proxy>
  authenticate (RADIUSRealm)
<Proxy>
  allow has_attribute.Service-Type=yes
  deny
```

- ❑ A group called **RegisteredUsersGroup** is allowed to access the ProxySG if the `allow group` gesture is defined.

```
<proxy>
  authenticate (RADIUSRealm)
<proxy>
  allow group=RegisteredUsersGroup
  deny
```

Troubleshooting

One of five conditions can cause the following error message:

Your request could not be processed because of a configuration error: "The request timed out while trying to authenticate. The authentication server may be busy or offline."

- ❑ The secret is wrong.
- ❑ The network is so busy that all packets were lost to the RADIUS server.
- ❑ The RADIUS server was slow enough that the ProxySG gave up before the server responded.
- ❑ The RADIUS servers are up, but the RADIUS server is not running. In this case, you might also receive ICMP messages that there is no listener.
- ❑ RADIUS servers machines are not running/unreachable. Depending on the network configuration, you might also receive ICMP messages.

Notes

If you use guest authentication, remember that RADIUS realms retrieve authorization data at the same time as the user is authenticated. In some cases, the system can distinguish between an authentication and authorization failure. Where the system cannot determine if the error was due to authentication or authorization, both the authentication and authorization are considered to be failed.

Chapter 59: Configuring the ProxySG as a Session Monitor

This chapter discusses how you can configure the SGOS software to monitor RADIUS accounting messages and to maintain a session table based on the information in these messages. The session table can then be used for logging or authentication.

You can also, optionally, configure multiple appliances to act as a session monitor *cluster*. When enabled, the session table is replicated to all members of the cluster to provide failover support.

After you configure and enable the session monitor, it maintains a session table that records which sessions are currently active and the user identity for each session. User information can be extracted from the session table by the ProxySG and used to make policy decisions.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ ["Configuring the Session Monitor"](#) on page 1099
- ❑ ["Session Monitor Attribute Substitutions"](#) on page 1104
- ❑ ["Creating the CPL"](#) on page 1105
- ❑ ["Access Logging"](#) on page 1105

Configuring the Session Monitor

To configure the session monitor, perform the following steps:

- ❑ Configure the RADIUS accounting protocol parameters for the session monitor.
- ❑ (Optional) Configure the session monitor cluster to handle failover.
- ❑ Configure the session monitor parameters.

Configuring the RADIUS Accounting Protocol Parameters

The configuration commands to create the RADIUS accounting protocol parameters can only be done through the CLI. If you are using session-monitor clustering, the commands must be invoked on each system in an already-existing failover group. (For information on configuring a failover group, see [Chapter 36: "Configuring Failover"](#) on page 807.)

To configure the RADIUS accounting protocol parameters:

- ❑ To enter configuration mode:

```
SGOS#(config) session-monitor
```


and

```
SGOS#(config) session-monitor attributes
```

- The following subcommands are available:

```
SGOS#(config session-monitor) radius acct-listen-port port_number
SGOS#(config session-monitor) radius authentication {enable | disable}
SGOS#(config session-monitor) radius encrypted-shared-secret
encrypted_secret
SGOS#(config session-monitor) radius no encrypted-shared-secret
SGOS#(config session-monitor) radius respond {enable | disable}
SGOS#(config session-monitor) radius shared-secret plaintext_secret
SGOS#(config session-monitor attributes) add attribute name | exit |
remove attribute name | view {calling-station-id | cisco-gateway-id}
```

Table 59–1 Session Monitor Accounting Command Descriptions

Command	Option	Description
<code>radius acct-listen-port</code>	<i>port_number</i>	The port number where the ProxySG listens for accounting messages
<code>radius authentication</code>	<code>enable</code> <code>disable</code>	Enable or disable (the default) the authentication of RADIUS messages using the shared secret. The shared secret must be configured before authentication is enabled.
<code>radius encrypted-shared-secret</code>	<i>encrypted_shared_secret</i>	Specify the shared secret (in encrypted form) used for RADIUS protocol authentication. The secret is decrypted using the configuration-passwords-key.
<code>radius no shared-secret</code>		Clears the shared secret used for RADIUS protocol authentication.
<code>radius respond</code>	<code>enable</code> <code>disable</code>	Enable (the default) or disable generation of RADIUS responses.
<code>radius shared-secret</code>	<i>plaintext_secret</i>	Specify the shared secret used for RADIUS protocol in plaintext.
<code>attributes</code>	add attribute <i>name</i> exit remove attribute <i>name</i> view { <i>calling-station-id</i> <i>cisco-gateway-id</i> }	Specify the RADIUS attributes that you want available as CPL substitutions, ELFF access log fields, and for authentication. The session monitor attributes must be identically defined under the RADIUS realm before they can be added under the session monitor. To define RADIUS realm attributes, see the Policy section in "RADIUS Realm Authentication and Authorization" on page 1089

Note: Any changes made to the Session-Monitor's attribute configuration will reinitialize the session table, resulting in the removal of all existing entries.

Note: To safely modify an existing user-defined attribute, you must first back up policy and remove the attribute from the RADIUS realm and session monitor configurations. See "[Example 5](#)" on page 1096 for instructions.

Configuring a Session Monitor Cluster

Configuring a session monitor cluster is optional. When a session monitor cluster is enabled, the session table is replicated to all members of the cluster. The cluster members are the ProxySG appliances that are configured as part of the failover group referenced in the session monitor cluster configuration. The failover group must be configured before the session monitor cluster. (For information on configuring a failover group, see [Chapter 36: "Configuring Failover"](#) on page 807.)

To replicate the session table to all the members of a failover group, you can use the following commands.

Note: When using a session monitor cluster, the RADIUS client must be configured to send the RADIUS accounting messages to the failover group's virtual IP address.

Proxy traffic can be routed to any of the machines in the cluster.

Note: Each member of the failover group must be *identically* configured to maintain the session table for RADIUS accounting messages.

To configure session monitor cluster parameters:

```
SGOS#(config) session-monitor
```

- ❑ The following subcommands are available:

```
SGOS#(config session-monitor) cluster {enable | disable}
```

```
SGOS#(config session-monitor) cluster group-address IP_address
```

```
SGOS#(config session-monitor) cluster port port_number
```

```
SGOS#(config session-monitor) cluster grace-period seconds
```

```
SGOS#(config session-monitor) cluster synchronization-delay seconds
```

```
SGOS#(config session-monitor) cluster retry-delay minutes
```

Table 59–2 Session Monitor Cluster Command Descriptions

Command	Option	Description
cluster	enable disable	Enable or disable (the default) clustering on a failover group. The group address must be set before the cluster can be enabled.
cluster group-address no group-address	IP_address	Set or clear (the default) the failover group IP address. This must be an existing failover group address.
cluster port	port_number	Set the TCP/IP port for the session replication control. The default is 55555.
cluster synchronization-delay	seconds	Set the maximum time to wait for session table synchronization. The default is zero; the range is from 0 to 2 ³¹ -1 seconds. During this time evaluation of \$(session-monitor.attribute) is delayed, so proxy traffic might also be delayed.

Table 59–2 Session Monitor Cluster Command Descriptions

Command	Option	Description
<code>cluster grace-period</code>	<i>seconds</i>	Set the time to keep session transactions in memory while waiting for slave logins. This can be set to allow session table synchronization to occur after the synchronization-delay has expired. The default is 30 seconds; the range is 0 to 2 ³¹ -1 seconds.
<code>cluster retry-delay</code>	<i>minutes</i>	Sets the maximum amount of time for connection retries in minutes. The delay can be set from 1 to 1,440 minutes.

Configuring the Session Monitor

The session monitor commands set up session monitoring behavior. If using session-monitor clustering, these commands must be invoked on all systems in the failover group.

To configure the session monitor:

1. At the (config) prompt:

```
SGOS#(config) session-monitor
SGOS#(config session-monitor) disable | enable
SGOS#(config session-monitor) max-entries integer
SGOS#(config session-monitor) timeout minutes
```

Table 59–3 Session Monitor Configuration Command Descriptions

Command	Option	Description
<code>enable disable</code>		Enable or disable (the default) session monitoring
<code>max_entries</code>	<i>integer</i>	The maximum number of entries in the session table. The default is 500,000; the range is from 1 to 2,000,000. If the table reaches the maximum, additional START messages are ignored.
<code>timeout</code>	<i>minutes</i>	The amount of time before a session table entry assumes a STOP message has been sent. The default is 120 minutes; the range is from 0 to 65535 minutes. Zero indicates no timeout.

2. (Optional) To view the session-monitor configuration, you can either use the `session-monitor view` command or the `config show session-monitor` command.

```
SGOS#(config) show session-monitor
General:
Status: enabled
Entry timeout: 120 minutes
Maximum entries: 500000
Cluster support: enabled
Cluster port: 55555
Cluster group address: 10.9.17.159
Synchronization delay: 0
Synchronization grace period: 30
```

```
Accounting protocol: radius
Radius accounting:
Listen ports:
Accounting: 1813
Responses: Enabled
Authentication: Enabled
Shared secret: *****
```

Session Monitor Attribute Substitutions

The attributes stored in the session table are available as CPL substitutions. These substitutions can be used to configure authentication within a valid policy substitution realm.

The session-monitor substitution uses the following syntax:

```
$(session-monitor.attribute.<attribute name>=)
```

Note: Session-monitor attribute names are *not* case-sensitive.

Testing Session Monitor CPL Attributes

The following CPL condition syntax can be used to test session-monitor CPL attributes:

```
session-monitor.attribute.<attribute name>=
```

The table below shows the supported comparison types for a given session-monitor attribute:

Table 59–4 Supported Attribute Comparison Methods

Attribute Type	Supported Comparisons
string	simple equality comparisons
integer	numerical range comparisons
IPv4/IPv6	IP address comparisons

All session-monitor attributes can use the following string comparison functions:

- *.prefix*
- *.suffix*
- *.substring*
- *.regex*

Attribute Comparison Examples

The following examples show the different types of attributes used in comparisons:

- ❑ **String:** `session-monitor.attribute.Calling-Station-ID="someuser"`
- ❑ **Integer:** `session-monitor.attribute.Framed-MTU=1`

- ❑ **IPv4:** `session-monitor.attribute.NAS-IP-Address=1.2.3.4`
- ❑ **IPv6:** `session-monitor.attribute.NAS-IPv6-Address=2001:db8:85a3::8a2e:370:7334`
- ❑ **Enum:** `session-monitor.attribute.Service-type=3`
`session-monitor.attribute.Service-type="Callback-Login"`

Note: The enum data type maps a string to an integer, and either can be used in comparisons. You can see a listing of the possible values for `Service-Type` (and other enum attributes) in the `security radius attributes` sub-mode.

Creating the CPL

Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file layers.

- ❑ In the following example, the ProxySG is using the session table maintained by the session monitor to extract user information for authentication.

```
<proxy>
  allow authenticate(session)
```

where `session` is a policy substitution realm that uses `$(session-monitor.radius.<attribute name>)` in building the username. (For information on creating a Policy Substitution realm, see [Chapter 57: "Policy Substitution Realm"](#) on page 1077.)

Access Logging

The Blue Coat ProxySG uses the following ELFF field syntax for access logging.

```
x-cs-session-monitor-radius(<attribute_name>)
```

When a user is authenticated by the ProxySG, the named attribute is fetched and recorded. When access log records are created, this field will be substituted with the value of the named attribute.

Access Logging is enabled on the **Configuration > Access Logging > General** page. For information about customizing access logs, see [Chapter 30: "Access Log Formats"](#) on page 657.

Notes

- ❑ The session table is stored entirely in memory. The amount of memory needed is roughly 40MB for 500,000 users.
- ❑ The session table is kept in memory. If the system goes down, the contents of the session table are lost. However, if the system is a member of a failover cluster, the current contents of the session table can be obtained from another machine in the cluster. The only situation in which the session table is entirely lost is if all machines in the cluster go down at the same time.
- ❑ The session replication protocol replicates session information only; configuration information is not exchanged. That means each ProxySG in the cluster must have identical RADIUS attribute settings in order to properly share information.
- ❑ The session replication protocol is not secured. The failover group should be on a physically secure network to communicate with each other.
- ❑ The session monitor requires sufficient memory and at least 100Mb-per-second network links among the cluster to manage large numbers of active sessions.
- ❑ The username in the session table is obtained from the Calling-Station-ID attribute in the RADIUS accounting message and can be a maximum of 19 bytes.

Chapter 60: Sequence Realm Authentication

This section describes how to configure the ProxySG to use multiple realms to authenticate users. It includes the following topics:

- ❑ ["About Sequencing"](#) on page 1107
- ❑ ["Adding Realms to a Sequence Realm"](#) on page 1107
- ❑ ["Creating a Sequence Realm"](#) on page 1108
- ❑ ["Defining Sequence Realm General Properties"](#) on page 1111
- ❑ ["Tips"](#) on page 1112

About Sequencing

After a realm is configured, you can associate it with other realms to allow the ProxySG to search for the proper authentication credentials for a specific user. That is, if the credentials are not acceptable to the first realm, they are sent to the second, and so on until a match is found or all the realms are exhausted. This is called *sequencing*.

For example, if a company has one set of end-users authenticating against an LDAP server and another using NTLM, a sequence realm can specify to attempt NTLM authentication first; if that fails because of a user-correctable error (such as credentials mismatch or a user not in database) then LDAP authentication can be specified to try next. You can also use sequences to fall through to a policy substitution realm if the user did not successfully authenticate against one of the earlier realms in the sequence.

Note: Errors such as *server down* do not fall through to the next realm in the sequence. Those errors result in an exception returned to the user. Only errors that are end-user correctable result in the next realm in the sequence being attempted.

Adding Realms to a Sequence Realm

Consider the following rules for using realm sequences:


- ❑ Ensure the realms to be added to the sequence are customized to your needs. Check each realm to be sure that the current values are correct. For IWA, verify that the **Allow Basic Credentials** option is set correctly.
- ❑ All realms in the realm sequence must exist and cannot be deleted or renamed while the realm sequence references them.
- ❑ Only one IWA realm is allowed in a realm sequence.
- ❑ If an IWA realm is in a realm sequence, it must be either the first or last realm in the list.

- ❑ If an IWA realm is in a realm sequence and the IWA realm does not support Basic credentials, the realm must be the first realm in the sequence and try IWA authentication once must be enabled.
- ❑ Multiple Basic realms are allowed.
- ❑ Multiple Windows SSO realms are allowed.
- ❑ Connection-based realms, such as Certificate, are not allowed in the realm sequence.
- ❑ A realm can only exist once in a particular realm sequence.
- ❑ A realm sequence cannot have another realm sequence as a member.
- ❑ If a realm is down, an exception page is returned. Authentication is not tried against the other later realms in the sequence.

Creating a Sequence Realm

To create a sequence realm:

1. Select the **Configuration > Authentication > Sequences > Sequence Realms** tab.
2. Click **New**. The Add Sequence Realm dialog displays.




3. In the **Realm name** field, enter a realm name. The name can be 32 characters long and composed of alphanumeric characters and underscores. The name must start with a letter.
4. Click **OK**.
5. Click **Apply**.


Adding Realms to a Sequence Realm

To add realms to a sequence realm:

1. Select the **Configuration > Authentication > Sequences > Sequence Main** tab.



2. Add a realm to the sequence:
 - a. Click **New**. The Member Realm dialog displays.
 - b. From the **Member Realm To Add** drop-down list, select an existing realm to the realm sequence. Remember that each realm can be used only once in a realm sequence.
 - c. Click **OK** to close the dialog.
3. To add additional realms to the sequence, repeat Step 2.
4. Click **Apply**.




5. To change the order that the realms are checked, use the **promote/demote** buttons. When you add an IWA realm, it is placed first in the list and you can allow the realm sequence to **try IWA authentication only once**. If you demote the IWA entry, it becomes last in the sequence and the default of checking IWA multiple times is enabled.
6. If you permit authentication or authorization errors, you can select the **Try next realm on tolerated error** checkbox to specify that the next realm on the list should be attempted if authentication in the previous realm has failed with a permitted error. The default value is to not attempt the next realm and fall out of the sequence. (For information on using permitted errors and guest authentication, see ["Permitting Users to Log in with Authentication or Authorization Failures"](#) on page 912.)
7. Click **Apply**.

Defining Sequence Realm General Properties

The **Sequence General** tab allows you to specify the display name and a virtual URL.

1. Select the **Configuration > Authentication > Sequences > Sequence General** tab.



The screenshot shows a web interface with three tabs: "Sequence Realms", "Sequence Main", and "Sequence General". The "Sequence General" tab is active. On the left, there are three numbered callouts (2, 3, 4) with arrows pointing to the corresponding fields. Field 2 is a drop-down menu for "Realm name" with "SequenceTest" selected. Field 3 is a text input for "Display name" with "SequenceTest" entered. Field 4 is a text input for "Virtual URL" with "www.cfauth.com/" entered. A mouse cursor is hovering over the Virtual URL field.

2. From the **Realm name** drop-down list, select the Sequence realm for which you want to change properties.
3. (Optional) If required, change the Sequence realm name in the **Display Name** field. The default value for the display name is the realm name. The display name cannot be longer than 128 characters and it cannot be null.
4. You can specify a virtual URL based on the individual realm sequence. For more information on the virtual URL, see "[Sequence Realm Authentication](#)" on page 1107.
5. Click **Apply**.

Tips

- ❑ Explicit Proxy involving a sequence realm configured with an NTLM/IWA realm and a substitution realm.

Internet Explorer (IE) automatically sends Windows credentials in the Proxy-Authorization: header when the ProxySG issues a challenge for NTLM/IWA. The prompt for username/password appears only if NTLM authentication fails. However, in the case of a sequence realm configured with an NTLM/IWA realm and a substitution realm, the client is authenticated as a guest in the policy substitution realm, and the prompt allowing the user to correct the NTLM credentials never appears.

- ❑ Transparent Proxy setup involving a sequence realm configured with an NTLM/IWA realm and a substitution realm.

The only way the ProxySG differentiates between a domain and non-domain user is through the NTLM/IWA credentials provided during the authentication challenge.

IE does not offer Windows credentials in the Proxy-Authorization: header when the Proxy issues a challenge for NTLM/IWA unless the browser is configured to do so. In this case, the behavior is the same as for explicit proxy.

If IE is not configured to offer Windows credentials, the browser issues a prompt for username/password, allowing non-domain users to be authenticated as guests in the policy substitution realm by entering worthless credentials.

Chapter 61: Managing X.509 Certificates

This section discusses X.509 certificates, which is a cryptographic standard for public key infrastructure (PKI) that specifies standard formats for public key certificates. Several RFCs and books exist on the public key cryptographic system (PKCS). This discussion of the elements of PKCS is relevant to their implementation in SGOS.

Blue Coat uses certificates for various applications, including:

- ❑ authenticating the identity of a server
- ❑ authenticating ProxySG
- ❑ securing an intranet
- ❑ encrypting data

Topics in this Section

This section includes the following topics:

- ❑ [Section A: "PKI Concepts"](#) on page 1114
- ❑ [Section B: "Using Keyrings and SSL Certificates"](#) on page 1118
- ❑ [Section C: "Managing Certificates"](#) on page 1130
- ❑ [Section D: "Using External Certificates"](#) on page 1138
- ❑ [Section E: "Advanced Configuration"](#) on page 1140
- ❑ [Section F: "Checking Certificate Revocation Status in Real Time \(OCSP\)"](#) on page 1151

Section A: PKI Concepts

The following sections describe the concepts of PKI (public key infrastructure) you must understand in order to use certificate authentication on the ProxySG appliance. The concepts included are the following:

- ❑ ["Public Keys and Private Keys"](#) on page 1114
- ❑ ["Certificates"](#) on page 1114
- ❑ ["Keyrings"](#) on page 1116
- ❑ ["Cipher Suites Supported by SGOS Software"](#) on page 1116
- ❑ ["Server-Gated Cryptography and International Step-Up"](#) on page 1117

Public Keys and Private Keys

In PKCS (public-key cryptography) systems, the intended recipient of encrypted data generates a private/public keypair, and publishes the public key, keeping the private key secret. The sender encrypts the data with the recipient's public key, and sends the encrypted data to the recipient. The recipient uses the corresponding private key to decrypt the data.

For two-way encrypted communication, the endpoints can exchange public keys, or one endpoint can choose a symmetric encryption key, encrypt it with the other endpoint's public key, and send it.

Certificates

Certificates are encrypted files that contain a public/private keypair. They can be used to verify the identity of a server, a website or to encrypt files.

The SGOS software uses:

- ❑ SSL Certificates.
- ❑ CA Certificates.
- ❑ External Certificates.
- ❑ Certificate Chains.

You can also use wildcard certificates during HTTPS termination. Microsoft's implementation of wildcard certificates is as described in RFC 2595, allowing an * (asterisk) in the leftmost-element of the server's common name only. For information on wildcards supported by Internet Explorer, refer to the Microsoft knowledge base, article: 258858. Any SSL certificate can contain a common name with wildcard characters.

SSL Certificates

SSL certificates are used to authenticate the identity of a server or a client. A certificate is confirmation of the association between an identity (expressed as a string of characters) and a public key. If a party can prove they hold the corresponding private key, you can conclude that the party is who the certificate says it is. The certificate contains other information, such as its expiration date.

The association between a public key and a particular server is done by generating a certificate signing request using the server's or client's public key. A certificate signing authority (CA) verifies the identity of the server or client and generates a signed certificate. The resulting certificate can then be offered by the server to clients (or from clients to servers) who can recognize the CA's signature. Such use of certificates issued by CAs has become the primary infrastructure for authentication of communications over the Internet.

The ProxySG appliance trusts all root CA certificates trusted by Internet Explorer and Firefox. The list is updated periodically to be in sync with the latest versions of IE and Firefox.

CA certificates installed on the ProxySG appliance are used to verify the certificates presented by HTTPS servers and the client certificates presented by browsers. Browsers offer a certificate if the server is configured to ask for one and an appropriate certificate is available to the browser.

For information on creating certificates, see ["Add Certificates to the ProxySG Appliance"](#) on page 1123

CA Certificates

CA certificates are certificates that belong to certificate authorities. CA certificates are used by ProxySG devices to verify X.509 certificates presented by a client or a server during secure communication.

ProxySG appliances are pre-installed with the most common CA certificates. You can review these certificates using the Management Console or the CLI. You can also add your own root and intermediate CA certificates for your own internal certificate authorities.

External Certificates

An external certificate is any X.509 certificate for which the ProxySG does not have the private key. The certificate can be used to encrypt data, such as access logs, with a public key so that it can only be decrypted by someone who has the corresponding private key. See ["Encrypting the Access Log"](#) on page 630 for information about encrypting access logs.

Certificate Chains

A certificate chain requires that certificates form a chain where the next certificate in the chain validates the previous certificate, going up the chain to the root, which is a trusted CA certificate.

Every certificate in the chain is checked for expiration as part of the certificate validation process. All certificates within this chain must be valid in order for the chain to be considered valid.

You can import certificate chains by creating a keyring and adding certificates to it. When creating certificate chains in the keyring, keep in mind that the keyring has a maximum character count of 7999. If you exceed the maximum, an error will appear on screen informing you that you have exceeded the character count limit.

In order for the ProxySG appliance to present a valid certificate chain for deployments such as **HTTPS SSL Forward Proxy** and **HTTPS Reverse Proxy**, the following measures must be taken:

- ❑ Add the server certificate to the keyring you created.

Following the server certificate, load any associated intermediate certificates in the certificate chain to the keyring. For detailed steps to create a certificate chain, see ["Importing a Server Certificate"](#) on page 1134.

Keyrings

A keyring contains a public/private keypair and can also contain a certificate signing request, a signed certificate and/or a certificate chain. Each keyring must have a name upon creation. You can view as well as delete a keyring. Some keyrings are already built-in for specified purposes. For information on managing keyrings, see ["Using Keyrings and SSL Certificates"](#) on page 1118.

Cipher Suites Supported by SGOS Software

A cipher suite specifies the algorithms used to secure an SSL connection. When a client makes an SSL connection to a server, it sends a list of the cipher suites that it supports.

The server compares this list with its own supported cipher suites and chooses the first cipher suite proposed by the client that they both support. Both the client and server then use this cipher suite to secure the connection.

Note: You can disable cipher suites that you do not trust. However, SGOS does not provide any mechanism to change the ordering of the ciphers used.

All cipher suites supported by the ProxySG use the RSA key exchange algorithm, which uses the public key encoded in the server's certificate to encrypt a piece of secret data for transfer from the client to server. This secret is then used at both endpoints to compute encryption keys.

By default, the ProxySG appliance is configured to allow TLSv1, TLSv1.1, and TLSv1.2 traffic. TECH24755 includes a list of cipher suites that are shipped with the appliance and that are available by default. The cipher suites available for use depend on the protocols you select.

Refer to the article on MySymantec:

<http://www.symantec.com/docs/TECH247556>

Note: Because they contain known vulnerabilities, Symantec recommends that you do not use the SSLv3 and SSLv2 protocols; however, if you do select the SSLv2 protocol, additional cipher suites are available: `DES-CBC3-MD5` (High, 168-bit), `RC2-CBC-MD5` (Medium, 128-bit), and `DES-CBC-MD5` (Low, 56-bit).

For information on cipher suite configuration, see ["Changing the Cipher Suite of the SSL Client"](#) on page 1167.

Note: Use of ECDHE ciphers is expected to increase CPU usage. On average, mid- and higher-end platforms could see up to a 10% increase in CPU utilization for typical traffic patterns. Lower-end platforms such as the 300 and 600, and older platforms such as the 810, may see a greater CPU increase. SGOS does not support ECDHE-ECDSA at this time. Support for ECDSA-based ciphers will be added in a future release.

Server-Gated Cryptography and International Step-Up

Due to US export restrictions, international access to a secure site requires that the site negotiates export-only ciphers. These are relatively weak ciphers ranging from 40-bit to 56-bit key lengths, and are vulnerable to attack.

Server Gated Cryptography (SGC) is a Microsoft extension to the certificate that allows the client receiving the certificate to first negotiate export strength ciphers, followed by a re-negotiation with strong ciphers. Netscape has a similar extension called International Step-up.

SGOS supports both SGC and International Step-up in its SSL implementation. There are, however, known anomalies in Internet Explorer's implementation that can cause SSL negotiation to fail. Refer to the following two documents for more detail and check for recent updates on the Microsoft support site.

<http://support.microsoft.com/support/kb/articles/Q249/8/63.ASP>
<http://support.microsoft.com/support/kb/articles/Q244/3/02.ASP>

To take advantage of this technology, SGOS supports VeriSign's Global ID Certificate product. The Global ID certificate contains the extra information necessary to implement SGC and International Step-up.

Section B: Using Keyrings and SSL Certificates

Keyrings are virtual containers. Each keyring holds a public/private key pair and a customized key length. You can associate certificates, certificate chains or certificate signing requests with keyrings.

In general, SSL certificates involve three parties:

- ❑ The subject of the certificate.
- ❑ The Certificate Authority (CA), which signs the certificate, attesting to the binding between the public key in the certificate and the subject.
- ❑ The *relying party*, which is the entity that trusts the CA and relies on the certificate to authenticate the subject.

Keyrings and certificates are used in:

- ❑ Encrypting data.
- ❑ Digitally Signing Access Logs.
- ❑ Authenticating end users.
- ❑ Authenticating a ProxySG appliance.

The steps in creating keyrings and certificates include:

- ❑ Create a keyring. A default keyring is shipped with the system and is used for accessing the Management Console, although you can use others. You can also use the default keyring for other purposes. You can create other keyrings for each SSL service. (See "[Creating a Keyring](#)" on page 1119.)

Note: You can also import keyrings. For information on importing keyrings, see "[Importing an Existing Keypair and Certificate](#)" on page 1140.

- ❑ (Optional) Create Certificate Signing Requests (CSRs) to be sent to Certificate Signing Authorities (CAs). (See "[Creating a CSR](#)" on page 1130.)
- ❑ Import X.509 certificates issued by trusted CA authorities for external use and associate them with the keyring. (See "[Managing SSL Certificates](#)" on page 1132.)

-or-

Create certificates and associate them with the keyring. (See "[Creating Self-Signed SSL Certificates](#)" on page 1133.)

Note: You can also associate a certificate chain with a keyring. For information on importing a certificate chain see, "[Importing a Server Certificate](#)" on page 1134

- ❑ (Optional, if using SSL Certificates from CAs) Import Certificate Revocation Lists (CRLs) so the ProxySG can verify that certificates are still valid.

Note: These steps must be done using a secure connection such as HTTPS, SSH, or a serial console.

Creating a Keyring

You can create additional keyrings for each HTTPS service defined.


The ProxySG appliance ships with several keyrings already created:

- ❑ **default:** The default keyring contains a certificate and an automatically-generated keyring and a self signed certificate which can be used for accessing the ProxySG appliance through HTTPS. As demonstrated by the ProxySG Management Console.
- ❑ **configuration-passwords-key:** The **configuration-passwords-key** keyring contains a keypair but does not contain a certificate. This keyring is used to encrypt passwords in the `show config` command and should not be used for other purposes.
- ❑ **appliance-key:** The **appliance-key** keyring contains an internally-generated keypair. If the ProxySG is authenticated (has obtained a certificate from the Blue Coat CA appliance-certificate server), that certificate is associated with this keyring, which is used to authenticate the device. (For more information on authenticating the ProxySG, see [Chapter 71: "Authenticating a ProxySG"](#) on page 1289.)
- ❑ **passive-attack-protection-only-key:** The **passive-attack-protection-only-key** keyring allows data to be encrypted, but with no endpoint authentication. Although the traffic cannot be sniffed, it can be intercepted with a man-in-the-middle attack. The **passive-attack-protection-only-key** keyring is NOT considered secure; therefore, it should not be used on production networks.


If an origin content server requires a client certificate and no keyring is associated with the ProxySG SSL client, the HTTPS connections fails. For information on using the SSL client, see [Chapter 62: "Managing SSL Traffic"](#) on page 1165.

To create a keyring:

1. Select the **Configuration > SSL > Keyrings > Keyrings** tab.



2. Click **Create**; the **Create Keyring** dialog displays.



3. Configure the options:

- a. **Keyring Name:** Give the keyring a meaningful name.

Note: Spaces in keyring names are not supported. Including a space can cause unexpected errors while using such keyrings.

- b. Select one of the following show options:

- **Show keypair** allows the keys to be viewed and exported.
- **Do not show keypair** prevents the keypair from being viewed or exported.
- **Show keypair to director** is a keyring viewable only if Director is issuing the command using a SSH-RSA connection.

Note: The choice among **show**, **do not show keypair**, and **show keypair to director** has implications for whether keyrings are included in profiles and backups created by Director. For more information, refer to the *Symantec Director Configuration and Management Guide*.

- c. Enter the key length in the **Create a new _____-bit keyring** field. The length range is 384-4096 bits. For deployments reaching outside the U.S., determine the maximum key length allowed for export.

Click **OK**. The keyring is created with the name you chose. It does not have a certificate associated with it yet. To associate a certificate or a certificate chain with a keyring, see "[Importing a Server Certificate](#)" on page 1134.

-or-

- d. Select **Import keyring**. The grayed-out **Keyring** field becomes enabled, allowing you to paste in an already existing private key. Any certificate or certificate request associated with this private key must be imported separately. For information on importing a certificate, see "[Importing a Server Certificate](#)" on page 1134.
- e. If the private key that is being imported has been encrypted with a password, select **Keyring Password** and enter the password into the field.

Note: The only way to retrieve a keyring's private key from the ProxySG is by using Director or the command line—it cannot be exported through the Management Console.

- 4. Click **OK** to close the dialog.
- 5. Click **Apply**.

To view or edit a keyring:

- 1. Select **Configuration > SSL > Keyrings > Keyrings**.
- 2. Click **Edit**.

Notes

- ❑ To view the keypair in an encrypted format, you can optionally specify `des` or `des3` before the `keyring_id`, along with an optional password. If the optional password is provided on the command line, the CLI does not prompt for a password.
- ❑ If the optional password is not provided on the command line, the CLI asks for the password (interactive). If you specify either `des` or `des3`, you are prompted.
- ❑ To view the keypair in unencrypted format, select either the optional `keyring_id` or use the `unencrypted` command option.
- ❑ You cannot view a keypair over a Telnet connection because of the risk that it could be intercepted.

Deleting an Existing Keyring and Certificate

To delete a keyring and the associated certificate:

1. Select the **Configuration > SSL > Keyrings > Keyrings** tab.
2. Highlight the name of the keyring to delete.
3. Click **Delete**. The Confirm delete dialog displays.
4. Click **OK** in the Confirm delete dialog.

Providing Client Certificates in Policy

Sometimes, when a user navigates to a secured Web address in a browser, the server hosting the site requests a certificate to authenticate the user. The client certificate authentication feature allows the ProxySG appliance to store client certificates and present the appropriate certificate to the Web server upon request. This feature is only applicable to intercepted SSL traffic.

The ProxySG appliance stores individual client certificates and keys in individual keyrings. You can then write policy that instructs the appliance which client certificate to use, and when to use it.

For convenience, you can also group client certificates and keyrings into a keylist that contains all of the client certificates for a specific purpose, such as certificates for a specific website or certificates for users in a particular group. If your policy references a keylist rather than an individual keyring, you must specify how to determine which certificate to use. This is done by matching the value of a substitution variable defined in the policy against a specified certificate field attribute value within the certificate. The ProxySG appliance determines what certificate field attribute to use based on an extractor string you supply when you create the keylist.

When a certificate is requested, if the policy selects a client certificate, the appliance presents the certificate to the requesting server. If no certificate is specified in policy, an empty certificate is presented.

Note: The ProxySG automatically detects and maintains a list of servers that request a client certificate during renegotiation. The appliance uses this list when evaluating the `client.certificate.requested` condition and correctly determines when a client certificate was requested during both the initial handshake and renegotiation. All additions to the list are event logged.

If the `client.certificate.requested` condition is removed from policy, no new entries are added to the list and the existing list remains unchanged until the condition is added again or the list is manually cleared.

To provide a client certificate to a requesting Web address, you must complete the following tasks.

Task #	Reference
1	"Add Certificates to the ProxySG Appliance" on page 1123
2	"Group Related Client Keyrings into a Keylist" on page 1125
3	"Specify the Client Certificates to be Used in Policy" on page 1127

Add Certificates to the ProxySG Appliance

Before certificates can be used in policy, they must be on the ProxySG appliance. Add the certificates to the appliance in one of the following ways:

- ["Create a Keydata File" on page 1124](#)

- ["Managing Certificates"](#) on page 1130

Create a Keydata File

Bundle multiple keyrings and keylists into a single keydata file for simple importing into the ProxySG appliance. The keydata file does not need to include both keyring and keylist information.

1. Open a new text file.
2. Add keyring information to the keydata file in the following format:

```
#keyring: <keyring_id>
#visibility: {show | show-director | no-show}
<Private Key>
<Certificate>
<CSR>
```

where:

- `keyring_id` - the name of the keyring.
- `visibility` - how the keyring is displayed in the `show configuration` output. Options include:
 - `show`: Private keys associated with keyrings created with this attribute can be displayed in the CLI or included as part of a profile or overlay pushed by Director.
 - `show-director`: Private keys associated with keyrings created with this attribute are part of the `show configuration` output if the CLI connection is secure (SSH/RSA) and the command is issued from Director.
 - `no-show`: Private keys associated with keyrings created with this attribute are not displayed in the `show configuration` output and cannot be part of a Director profile. The `no-show` option is provided as additional security for environments where the keys will never be used outside of the particular ProxySG.
- `Private Key, Certificate, and CSR` - Paste the contents of the key, certificate or CSR into the text file, including the `---Begin` and `---End` tags.

In the following example, the private key and certificate has been truncated.

```
#keyring:Keyring1
#visibility:no-show
-----BEGIN RSA PRIVATE KEY-----
MIIEpAIBAAKCAQE...KvBgDmSIw6dTXxAT/mMUHGRd7cRew==
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
MIIDdjCCA14CCQC...TjUwxwboMEyL60z/tixM=
-----END CERTIFICATE-----
#keyring:Keyring2
```

3. Add keylist information to the file in the following format:

```
#keylist: <keylist_name>
#extractor: <extractor>
<keyring_id>
<keyring_id>
```

where:

- `keylist_name` - Type the name of the keylist.
- `extractor` - Enter a string that identifies which certificate field attribute value to extract to determine a policy match, using the `$(field.attribute)` syntax. Substitutions from all attributes of Subject, Issuer, SubjectAltName, IssuerAltName, and SerialNumber certificate fields are supported.
- `keyring_id` - List any keyrings to include in the keylist. The keyrings may be included in the keydata file, or may be keyrings that already exist on the ProxySG appliance.

For example:

```
#keylist:mylist
#extractor: $(Subject.CN)
Keyring1
Keyring2
```

4. Save the file as .txt on a web server that can be accessed by the ProxySG appliance.

Import Certificates onto the ProxySG Appliance

Use the following procedure to import multiple client certificates (as well as the associated key pair and CSR) into the ProxySG appliance.

1. Select **Configuration > SSL > Keyrings > Import**.
2. In the **URL** field type the path to the keydata file with the keylists and keyrings.
3. (Optional) If you have encrypted the private keys in the keydata file, type the Passphrase for the private keys.

All keyrings or keylists being imported must have the same Passphrase for the import to be successful.

4. Click **Import**.
5. Click **Apply**.

Group Related Client Keyrings into a Keylist

To easily reference client certificate keyrings in policy, use keylists to group them together. For example it is often useful to group certificates into keylists bundled by:

- all client certificates for a specific web address,
- all client certificates for a group of users,
- all client certificates for a specific user.

All keyrings in the keylist must have the same extractor, but each certificate must have a unique value for the extractor. The evaluation of the keylist extractor string must be unique across the client certificates in the keylist, otherwise changes being applied to the keylist will fail with an error.

1. Select **Configuration > SSL > Keyrings > Keylists**.
2. Click **Create**.
3. In the **Name** field, type a name for the new keylist.
4. In the **Extractor** field enter a string that identifies which certificate field attribute value to extract to determine a policy match. Enter the string using the `$(field.attribute)` syntax. For example, to extract the value of the CN attribute from the Subject field of the certificate, you would enter `$(subject.CN)`.

Alternatively, select values from the **Field**, **Attribute**, and **Group Name** drop down lists to build an extractor string, and click **Add to extractor**. The new extractor string is appended to any existing text in the **Extractor** field. The Group Name drop down list only appears for IssuerAltName and SubjectAltName fields. The Extractor field can have a maximum of 255 characters.

The extractor supports substitutions from all attributes of Subject, Issuer, SubjectAltName, IssuerAltName, and SerialNumber certificate fields. The default extractor value is `$(Subject.CN)`; many other subject attributes are recognized, among them OU, O, L, ST, C, and DC. Field indexes can be used in substitutions on a group name or attribute; for example `$(SubjectAltName.DNS.1)`.

5. From the **Available Keyrings** list, select the keyrings to be included in this keylist and click **Add**.

To remove a keyring from the list of **Included Keyrings**, select the keyring and click **Remove**.

If any errors are noted in the Included Keyrings list, the keylist cannot be created. Possible causes for errors are:

- The included keyring does not contain the specified extractor pattern or substitution variable.
- Multiple keyrings have the same value for the specified extractor.

The extracted value in the keyring allows the policy action object to find the appropriate keyring certificate to use. Only one keyring can be utilized by each policy transaction. Therefore, the extractor string evaluation must be unique across the certificates in the keylist. A keyring whose extractor value matches the extractor value of any existing keyring in the keylist will not be added to the keylist. For example, if the extractor `$(Subject.DC)` is selected, and all keyrings have the same value in the certificate for that extractor, the policy would not be able to determine which keyring to select.

6. To save the keylist click **OK**.

7. Click **Apply**.


Specify the Client Certificates to be Used in Policy

You can now reference the keyrings and keylists in your policy.


Specify the Client Certificates to be Used in Policy in the VPM

To respond to client certificate requests, in the SSL Access policy layer add an action object with the keyrings or keylists that can provide client certificates when requested.

To use a keyring



1. In the **Name** field, enter a name for the object or leave as is to accept the default.
2. Select **Keyring**.
3. From the drop-down, select the keyring to use in policy.
4. Click **OK**.

To use a keylist

1. In the **Name** field, enter a name for the object or leave as is to accept the default.
2. Select **Keylist**.
3. From the drop-down, select the keylist to use in policy.
4. In the **Selector** field, type a substitution variable.

All substitution variables are supported; however recommended substitution variables for the selector include `$(user)`, `$(group)`, and `$(server.address)`. For information on substitution variables, see "CPL Substitutions" on page 495 of the *Content Policy Language Reference*.

Note: The Selector value must match the set of extractor values that are displayed when you run the `view` command for a keylist. For example, if the Subject.CN in the certificate is set to represent a user name, use the Selector `$(user)`, and select the Extractor value `$(Subject.CN)`. If the Extractor value was set to `$(Subject.O)`, no match would be found and a certificate would not be sent.

If you are using the `$(group)` selector, you must also create a list of the groups to be included in the `$(group)` substitution variable. See "Creating the Group Log Order List" in the *Visual Policy Manager Reference*.

5. Click **OK**.

Specify the Client Certificates to be Used in Policy in CPL

To respond to client certificate requests, add a keyring or keylist with the following syntax in the `<SSL>` layer:

```
server.connection.client_keyring(keyring)
server.connection.client_keyring(keylist, selector)
```

where:

- *keyring*—Specifies the keyring to use for client certificate requests.

- *keylist*—Specifies the keylist to use for client certificate requests. The *selector* value must also be specified.
- *selector*—Takes a substitution variable.

All substitution variables are supported; however recommended substitution variables for the selector include `$(user)`, `$(group)`, and `$(server.address)`.

Keyring Examples

- ❑ Use the certificate from `<keyring>` as the client certificate for user `<user>` connecting to a specific website `<url>`.

```
url=<url> user=<user> server.connection.client_keyring(<keyring>)
```
- ❑ Use the certificate from `<keyring>` as the client certificate for user `<user>` connecting to any website that requires a client certificate.

```
user=<user> server.connection.client_keyring(<keyring>)
```
- ❑ Use the certificate from `<keyring>` as the client certificate for all users of group `<group>` connecting to a specific website `<url>`.

```
url=<url> group=<group> server.connection.client_keyring(<keyring>)
```

Keylist Examples

- ❑ Select a keyring or certificate from the keylist `<keylist>` whose extractor value is equal to the user of the connection, for a specific website `<url>`.

```
<SSL>
url = <url> server.connection.client_keyring(<keylist>, \
    "$(user)")
```
- ❑ For connections to a website `<url>`, this will select a keyring or certificate from keylist `<keylist>` whose extractor value is equal to the group of the connection.

```
<SSL>
url = <url> group = (<group>, <group>) \
    server.connection.client_keyring(<keylist>, "$(group)")
```

Section C: Managing Certificates

This section discusses how to manage certificates, from obtaining certificate signing requests to using certificate revocation lists.

In this section are:

- ❑ "Managing Certificate Signing Requests" on page 1130
- ❑ "Managing SSL Certificates" on page 1132
- ❑ "Using Certificate Revocation Lists" on page 1134
- ❑ "Troubleshooting Certificate Problems" on page 1136


Managing Certificate Signing Requests

Certificate signing requests (CSRs) are used to obtain a certificate signed by a Certificate Authority. You can also create CSRs off box.


Creating a CSR

To create a CSR:


1. Select the **Configuration > SSL > Keyrings** tab.



2. Select the keyring for which you need a signed certificate and click **Edit**. The Edit Keyring dialog displays.



3. In the **Certificate Signing Request** area, click **Create**. The Create Certificate-signing-request dialog displays.



4. Fill in the fields:
 - **State/Province**—Enter the state or province where the machine is located.
 - **Country Code**—Enter the two-character ISO code of the country.
 - **City/Locality**—Enter the city.
 - **Organization**—Enter the name of the company.
 - **Unit**—Enter the name of the group that is managing the machine.
 - **Common Name**—Enter the URL of the company.
 - **Challenge**—Enter a 4-20 character alphanumeric challenge.

- **E-mail Address**—The e-mail address you enter must be 60 characters or less. A longer e-mail address generates an error.
- **Company**—Enter the name of the company.

Note: Most field limits are counted in terms of bytes rather than characters. The number of non-ASCII characters a field can accommodate will be less than the size limit because non-ASCII characters can occupy more than one byte, depending on the encoding. The only exception is the **Challenge** field, which is counted in terms of characters.

5. Click **OK** to close the dialog. The **Certificate Signing Request** area displays the certificate information.
6. Click **OK** to close the dialog. The **CSR** column for the keyring displays **Yes**.

Viewing a Certificate Signing Request

After a CSR is created, you must submit it to a CA in the format the CA requires. You can view the output of a certificate signing request.

To view the output of a certificate signing request:

1. Select the **Configuration > SSL > Keyrings** tab.
2. Click **Edit**.
3. From the drop-down list, select the keyring for which you have created a certificate signing request.

The certificate signing request displays in the Certificate Signing Request window and can be copied for submission to a CA.

Managing SSL Certificates

SSL certificates can be obtained two ways:

- ❑ Created on the ProxySG as a self-signed certificate

To create a SSL self-signed certificate on the ProxySG using a Certificate Signing Request, continue with the next section.

- ❑ Imported after receiving the certificate from the signing authority.

If you plan to use SSL certificates issued by Certificate Authorities, the procedure is:

- Obtain the keypair and Certificate Signing Requests (CSRs), either off box or on box, and send them to the Certificate Authority for signing.
- After the signed request is returned to you from the CA, you can import the certificate into the ProxySG. To import a certificate, see "[Importing a Server Certificate](#)" on page 1134.

Creating Self-Signed SSL Certificates

The ProxySG appliance ships with a self-signed certificate, which is associated with the default keyring. Only one certificate can be associated with a keyring. If you have multiple uses, use a different keyring and associated certificate for each one. Self-signed certificates are generally meant for intranet use, not Internet.

To create a self-signed certificate:

1. Select the **Configuration > SSL > Keyrings > Keyrings** tab.
2. Highlight the keyring for which you want to add a certificate.
3. Click **Edit** in the **Keyring** tab.
4. Click **Create**.

5. Fill in the fields:
 - **State/Province**—Enter the state or province where the machine is located.
 - **Country Code**—Enter the two-character ISO code of the country.
 - **City/Locality**—Enter the city.
 - **Organization**—Enter the name of the company.
 - **Unit**—Enter the name of the group that is managing the machine.
 - **Common Name**—A common name should be the one that contains the URL with client access to that particular origin server.
 - **Challenge**—Enter a 4-20 character alphanumeric challenge.
 - **E-mail Address**—The e-mail address you enter must be 60 characters or less. A longer e-mail address generates an error.
 - **Company**—Enter the name of the company.

The **Create** tab displays the message: **Creating.....**

Importing a Server Certificate

Once your certificate is approved by the signing authority, you can import your server certificate onto the ProxySG appliance and associate it with a keyring. You can also import a certificate chain to be associated with a keyring as detailed in the steps below.

To import a server certificate:

The steps below will also guide you through importing a certificate chain. Certificate chains require that you import your server certificate first followed by all associated intermediate certificates.

1. Copy the certificate to your clipboard. You must include the “-----BEGIN CERTIFICATE-----” and “-----END CERTIFICATE-----” statements.
2. Select **Configuration > SSL > Keyrings**.
3. Highlight the keyring for which you want to import a certificate.
4. Click **Edit** in the **Keyrings** tab.
5. In the **Certificate** panel, click **Import**.
6. Paste the certificate you copied into the dialog box.
7. For certificate chains, copy each intermediate certificate to your clipboard individually then paste each certificate you copied in to the **Import Certificate** dialog. You must include the “-----BEGIN CERTIFICATE-----” and “-----END CERTIFICATE-----” statements. Intermediate certificates must follow the server certificate.

Repeat step 7 until you have copied and pasted all associated intermediate certificates.

Note: Certificate chains, when imported to a keyring, have a maximum character count of 7999. If you exceed the maximum character count, the management console will inform you by displaying an error message on your screen.

8. Click **OK**.
9. Click **Apply**.

The SSL Certificate Pane, displays the certificate(s) and it's associated keyring.

Using Certificate Revocation Lists

Certificate Revocation Lists (CRLs) enable checking server and client certificates against lists provided and maintained by CAs that show certificates that are no longer valid. Only CRLs that are issued by a trusted issuer can be successfully verified by the ProxySG. The CRL can be imported only when the CRL issuer certificate exists as a CA certificate on the ProxySG.

You can determine if the ProxySG SSL certificates are still valid by checking *Certificate Revocation Lists* (CRLs) that are created and issued by trusted Certificate Signing Authorities. A certificate on the list is no longer valid.

Only CRLs that are issued by a trusted issuer can be verified by the ProxySG successfully. The CRL can be imported only when the CRL issuer certificate exists as a CA certificate on the ProxySG.

SGOS allows:

- ❑ One local CRL list per certificate issuing authority.
- ❑ An import of a CRL that is expired; a warning is displayed in the log.
- ❑ An import of a CRL that is effective in the future; a warning is displayed in the log.

CRLs can be used for the following purposes:

- ❑ Checking revocation status of client or server certificates with HTTPS Reverse Proxy.
- ❑ Checking revocation status of client or server certificates with SSL proxy. (For more information on using CRLs with the SSL proxy, see "[Validating the Server Certificate](#)" on page 214.)
- ❑ ProxySG-originated HTTPS downloads (secure image download, content filter database download, and the like).
- ❑ PEM-encoded CRLs, if cut and pasted through the inline command. Refer to the *Command Line Interface Reference* for more information.
- ❑ DER-format (binary) CRLs, if downloaded from a URL.

To import a CRL:

You can choose from among four methods to install a CRL on the ProxySG:

- ❑ Use the Text Editor, which allows you to enter the installable list (or copy and paste the contents of an already-created file) directly onto the ProxySG appliance.
- ❑ Create a local file on your local system.
- ❑ Enter a remote URL, where you placed an already-created file on an FTP or HTTP server to be downloaded to the ProxySG.
- ❑ Use the CLI `inline` command. Refer to the *Command Line Interface Reference* for more information.

To update a CRL:

1. Select the **Configuration > SSL > CRLs** tab.
2. Click **New** or highlight an existing CRL and click **Edit**.
3. Give the CRL a name.
4. From the drop-down list, select the method to use to install the CRL; click **Install**.

- Remote URL:
Enter the fully-qualified URL, including the filename, where the CRL is located. To view the file before installing it, click **View**. Click **Install**.
The **Install CRL** dialog displays. Examine the installation status that displays; click **OK**.
- Local File:
Click **Browse** to display the Local File Browse window. Browse for the CRL file on the local system. Open it and click **Install**. When the installation is complete, a results window opens. View the results, close the window, click **Close**.
- Text Editor:
Copy a new CRL file into the window, and click **Install**.
When the installation is complete, a results window opens. View the results, close the window, click **Close**.

Note: The Management Console text editor can be used to enter a CRL file. You cannot use it to enter CLI commands.

5. Click **OK**; click **Apply**

Troubleshooting Certificate Problems

Two common certificate problems are discussed below.

- ❑ If the client does not trust the Certificate Signing Authority that has signed the appliance's certificate, an error message similar to the following appears in the event log:

```
2004-02-13 07:29:28-05:00EST "CFSSL:SSL_accept error:14094416:SSL
routines:SSL3_READ_BYTES:ssl3 alert certificate unknown" 0
310000:1
../cf_ssl.cpp:1398
```

This commonly occurs when you use the HTTPS-Console service on port 8082, which uses a self-signed certificate by default. When you access the Management Console over HTTPS, the browser displays a pop-up that says that the security certificate is not trusted and asks if you want to proceed. If you select **No** instead of proceeding, the browser sends an *unknown CA alert* to the ProxySG.

You can eliminate the error message one of two ways:

- If this was caused by the Blue Coat self-signed certificate (the certificate associated with the default keyring), import the certificate as a trusted Certificate Signing Authority certificate. See ["Importing a Server Certificate"](#) on page 1134 for more information.
- Import a certificate on the ProxySG for use with HTTPS-Console that is signed by a CA that a browser already trusts.

- ❑ If the ProxySG appliance's certificate is not accepted because of a *host name mismatch* or it is an *invalid certificate*, you can correct the problem by creating a new certificate and editing the HTTPS-Console service to use it. For information on editing the HTTPS-Console service, see "[Managing the HTTPS Console \(Secure Console\)](#)" on page 1270.


Section D: Using External Certificates

External certificates are certificates for which Blue Coat does not have the private key. The first step in using external certificates is to import the certificates onto the ProxySG appliance.

Importing and Deleting External Certificates

To Import an external certificate:

1. Copy the certificate onto the clipboard.
2. Select the **Configuration > SSL > External Certificates** tab.
3. Click **Import**.



4. Enter the name of the external certificate into the **External Cert Name** field and paste the certificate into the **External Certificate** field. You must include the "-----BEGIN CERTIFICATE-----" and "-----END CERTIFICATE-----" statements.
5. Click **OK**.
6. Click **Apply**.

Deleting an External Certificate

To delete an external certificate:

1. Select the **Configuration > SSL > External Certificates** tab.

2. Highlight the name of the external certificate to be deleted.
3. Click **Delete**.
4. Click **OK** in the Confirm Delete dialog that displays.
5. Click **Apply**.

Digitally Signing Access Logs

You can digitally sign access logs to certify that a particular ProxySG wrote and uploaded a specific log file. Signing is supported for both content types—text and gzip—and for both upload types—continuous and periodic. Each log file has a signature file associated with it that contains the certificate and the digital signature used for verifying the log file. When you create a signing keyring (which must be done before you enable digital signing), keep in mind the following:

- ❑ The keyring must include a certificate. .
- ❑ The certificate purpose must be set for **smime** signing. If the certificate purpose is set to anything else, you cannot use the certificate for signing.
- ❑ Add the `%c` parameter in the filenames format string to identify the keyring used for signing. If encryption is enabled along with signing, the `%c` parameter expands to `keyringName_Certname`.

For more information about digitally signing access logs, see "[Encrypting the Access Log](#)" on page 630.

Section E: Advanced Configuration

This section includes the following topics:

- ❑ "Importing an Existing Keypair and Certificate" on page 1140
- ❑ "Importing CA Certificates" on page 1141
- ❑ "Managing CA Certificate Lists" on page 1143

Importing an Existing Keypair and Certificate

If you have a keypair and certificate used on one system, you can import that same keypair and certificate for use on a different system. You can also import a certificate chain. Use the `inline certificate` command to import multiple certificates through the CLI. Refer to the *Command Line Interface Reference* for more information.

If you are importing a keyring and one or more certificates onto a ProxySG, first import the keyring, followed by the related certificates. The certificates contain the public key from the keyring, and the keyring and certificates are related.

To Import a keyring:

1. Copy the already-created keypair onto the clipboard.
2. Select the **Configuration > SSL > Keyrings > SSL Keyrings** tab.
3. If the keyring already exists, select the keyring and click **Delete** and **Apply**.
4. Click **Create**. The Create Keyring dialog displays.

5. Configure the keyring options:
 - a. Select a show option:
 - **Show keypair** allows the keys to be exported.
 - **Do not show keypair** prevents the keypair from being exported.

- **Show keypair to director** is a keyring viewable only if Director is issuing the command using a SSH-RSA connection.

Note: The choice among **show**, **do not show** and **show keypair to director** has implications for whether keyrings are included in profiles and backups created by Director. For more information, refer to the *Blue Coat Director Configuration and Management Guide*.

- b. Select the **Import keyring** option.

The grayed-out **Keyring** field becomes enabled, allowing you to paste in the already existing keypair. The certificate associated with this keypair must be imported separately.


- c. If the keypair that is being imported has been encrypted with a password, select **Keyring Password** and enter the password into the field.
- d. Click **OK**.

6. Click **Apply**.

Importing CA Certificates

The ProxySG appliance is preinstalled with and trusts all root CA certificates trusted by Internet Explorer and Firefox. This certificate list is updated periodically to be in sync with the latest versions of IE and Firefox.

You can also import non-standard third party CA certificates into the ProxySG CA certificate store, including root and intermediate CA certificates. By adding CA certificates to the CA certificate store, these will be available for use by the CA certificate lists (CCL) for validating the security of connections.



To Import a CA Certificate:

1. Click **Import**. The Import CA Certificate dialog displays.
2. Name the certificate.

Note: Spaces in CA Certificate names are *not* supported. Including a space can cause unexpected errors while using such certificates.

3. Paste the signed CA Certificate into the **Import CA Certificate** field.
4. Click **OK**.
5. Click **Apply**.

To Import a CA Certificate and associate it with a keyring:

1. Copy the certificate onto the clipboard.
2. Select **Configuration > SSL > Keyrings** and click **Edit/View**.
3. From the drop-down list, select the keyring that you just imported.
4. Click **Import** in the Certificate field.
5. Paste the certificate into the **Import Certificate** dialog that appears. Be sure to include the "-----BEGIN CERTIFICATE-----" and "-----END CERTIFICATE-----" statements.

6. For certificate chains, repeat step 5. You must copy and paste each associated intermediate certificate individually into the keyring.

Note: Certificate chains, when imported into a keyring, have a maximum character count of 7999. If you exceed the maximum character count, the management console inform you by displaying an error message on your screen.

7. Click **OK**.
8. Click **Apply**.

To view a CA certificate:

1. Select the **Configuration > SSL > CA Certificates > CA Certificates** tab.
2. Select the certificate you want to view.
3. Click **View**. Examine the contents and click **Close**.

To delete a CA certificate:

1. Select the **Configuration > SSL > CA Certificates > CA Certificates** tab.
2. Select the certificate to delete.
3. Click **Delete**.
4. Click **OK**.

Managing CA Certificate Lists

A CA certificate list (CCL), which contains some of the CA Certificates available on the ProxySG, allows the administrator to control the set of CA certificates trusted for a particular set of SSL connections. A CCL contains a subset of the available CA certificates on the ProxySG, and restricts trust to those certificates. The CCL referenced by the profile or service configuration is used when an SSL connection is established to that service or using that profile.

Three CCLs are created by default on the ProxySG appliance:

- ❑ `appliance-ccl`: This CCL is used for authenticating connections among devices manufactured by Blue Coat Systems. By default it contains the Blue Coat ABRCA root certificate (ABRCA_root).

This list is used by default in the **bluecoat-appliance-certificate** SSL device profile. This CCL can be edited but not deleted.

For more information on device authentication, see [Chapter 71: "Authenticating a ProxySG"](#) on page 1289.

- ❑ `browser-trusted`: This CCL includes most of the well-known CAs trusted by common browsers such as Internet Explorer and Firefox. This CCL can be edited but not deleted. You can manually add CAs to this list. In addition, the ProxySG appliance automatically retrieves an updated `browser-trusted` CCL from Symantec every seven days. For information on how to customize the automatic update behavior, see ["Configure Automatic Updates"](#) on page 1147. The `browser-trusted` CCL is used by default during certificate verification by the SSL client and by the **default** SSL device profile.
- ❑ `image-validation`: This CCL is used to validate signed SGOS images.


Note: For information on using the SSL client or SSL device profiles, see [Chapter 62: "Managing SSL Traffic"](#) on page 1165.

You can customize the CCLs available on the ProxySG appliance to ensure that the appliance has the CA certificates it needs to handle HTTPS requests. You can create your own CA certificate lists or modify the default CCLs by adding or removing trusted CAs:

- ❑ ["Creating a CA Certificate List:"](#) on page 1144
- ❑ ["Updating a CA Certificate List"](#) on page 1146
- ❑ ["Configuring Download of CCL Updates from Symantec"](#) on page 1146

Creating a CA Certificate List:

1. Select **Configuration > SSL > CA Certificates > CA Certificate Lists**.



2. Configure the list:
 - a. Click **New** to create a new list. The Create CA Certificate List dialog displays.
 - b. Enter a meaningful name for the list in the **CA-Certificate List Name** field.
 - c. Add or remove CAs from the list as follows:
 - To add CA Certificates to the list, highlight the certificate and click **Add**. The certificate *must* have been imported onto the ProxySG appliance before it can be added to a certificate list. See ["Importing CA Certificates"](#) on page 1141.
 - To remove CA Certificates from the list, highlight the certificate in the **Add** list and click **Remove**.
 - d. Click **OK**
3. Click **Apply**.

Updating a CA Certificate List

Because the list of trusted CAs changes over time, you may want to update your CCLs to ensure that they contain the most up-to-date list of CA certificates. You can manually edit the default `appliance-ccl` and `browser-trusted` CCLs as well as any custom-produced CCL. The `bluecoat-services` and `image-validation` CCLs are read-only and cannot be modified by the user; however, you can still view the contents.

Keep in mind that if you plan to add a CA to a CCL, you must first import the corresponding CA certificate as described in ["Importing CA Certificates"](#) on page 1141.

For the `browser-trusted` CCL, you also have the option to configure the appliance to download an updated browser-trusted list of CAs on demand or automatically on a schedule. This smart download compares the existing browser-trusted list on the appliance to the new list and only adds those CA certificates that are new since the last update. Any manual changes that you have made to the file are preserved.

To update a CCL manually

1. Select **Configuration > SSL > CA Certificates > CA Certificate Lists**.
2. Select the CCL you want to modify and click **Edit**.
 - a. Add or remove CAs from the list as follows:
 - To add CA Certificates to the list, highlight the certificate and click **Add**. The certificate *must* have been imported onto the ProxySG appliance before it can be added to a certificate list. See ["Importing CA Certificates"](#) on page 1141.
 - To remove CA Certificates from the list, highlight the certificate in the **Add** list and click **Remove**.
 - b. Click **OK**.
3. Click **Apply**.

Configuring Download of CCL Updates from Symantec

By default, the ProxySG appliance will automatically download and install a package containing the updated CA Certificates and CCL updates—called a *trust package*—from the Symantec website every seven days. This trust package contains any updates to the `browser-trusted` and `image-validation` CCLs and their associated CA certificates since the last update, based on the timestamp at the time the trust package was created. Note that any manual changes you have made to the CCLs and CA certificates will be preserved.

You can customize the CA download list updates as follows:

- ❑ ["Change the Download Location"](#) on page 1147
- ❑ ["Configure Automatic Updates"](#) on page 1147
- ❑ ["Load the Trust Package"](#) on page 1148

- "Verify Trust Package Downloads" on page 1148

Change the Download Location

The downloadable CA list—called a *trust package*—is hosted at the following URL:

http://appliance.bluecoat.com/sgos/trust_package.bctp

By default, the ProxySG appliance is configured to download the trust package directly from this URL. As an alternative you can set up your own download site on premise. To do this, you must download the trust package from the URL to your download server and then configure the download path on the appliances in your network.

After you determine the download location, you must configure the appliance to point to the location using the following command:

```
#(config) security trust-package download-path <URL>
```

For example, to configure the appliance to download the trust package from a `bluecoat` folder on your `download.acme.com` server, you would enter the following command:

```
#(config) security trust-package download-path http://  
downloads.acme.com/bluecoat/trust_package.bctp
```

Note: The ProxySG appliance can only download and install a `trust_package.bctp` trust package created by Symantec.

Configure Automatic Updates

By default, the ProxySG appliance automatically downloads and installs the latest trust package every seven days by default. You can disable automatic updates or modify the update interval as follows:

To disable automatic updates

If you prefer to manually download and install the trust package, you can download automatic updates as follows:

```
#(config) security trust-package auto-update disable
```

To change the update interval

```
#(config) security trust-package auto-update interval <days>
```

where `<days>` is the number of days between updates. This value can be from 1 to 30 inclusive. For example, to set the auto-update interval to 10 days, you would enter the following command:

```
#(config) security trust-package auto-update interval 10
```

To enable automatic updates

If you previously disabled automatic updates, you can re-enable them using the following command:

```
#(config) security trust-package auto-update enable
```

Note that if you previously modified the automatic update interval, your settings will be preserved.

Load the Trust Package

If you want to manually download and install the trust package—either because you have disabled automatic updates or you want to force an update before the next automatic update—enter the following command:

```
#load trust-package
Downloading from "http://appliance.bluecoat.com/sgos/
trust_package.bctp"
    The trust package has been successfully downloaded.
    trust package successfully installed
```

Verify Trust Package Downloads

Use the following command to view the status of the trust package downloads:

```
#show security trust-package
Download url: http://appliance.bluecoat.com/sgos/trust_package.bctp
Auto-update: enabled           Auto-update interval: 7 days

Previous (success) install via manual

Creation time: Saturday October 1 2011 00:26:43 UTC

CA Certificate List changes:
    browser-trusted: CAs - 3 added, 4 deleted, 0 modified

image-validation install: Tuesday October 11 2011 00:26:27 UTC

Download log:
    Downloaded at: Tuesday October 11 2011 00:26:27 Success

    Downloaded from: http://appliance.bluecoat.com/sgos/
    trust_package.bctp
```

Managing Cached Intermediate Certificates

The ProxySG appliance automatically stores unrecognized intermediate CA certificates that are included with validated CA certificate chains whenever an SSL connection is established.

These intermediate CA certificates are stored within a separate cache on the appliance and are used to validate SSL connections when an incomplete certificate chain is encountered. For security purposes, OCSP and CRL validation checks are performed to confirm the safety of the certificate chain. As an additional layer of security, the intermediate CA certificates in the chain must end

with a trusted root certificate from the CCL (CA certificate list) that is associated with the connection. If a compatible certificate is not found, the connection is considered insecure and the user will be given a security warning.

Note: The ProxySG appliance does not allow automatic retrieval of issuing certificates for Intermediate certificates that include an AIA (Authority Information Access) entry.

You can control the following aspects of Intermediate Certificate Caching:

- ❑ ["Turn off Intermediate Certificate Caching"](#) on page 1149
- ❑ ["View Cached Intermediate Certificates"](#) on page 1149
- ❑ ["Clear Cached Intermediate Certificates"](#) on page 1150

Turn off Intermediate Certificate Caching

Turning off caching automatically clears the existing cache of intermediate CA certificates and prevents any validated intermediate certificates from being added to the cache.

To turn off intermediate certificate caching:

1. Select **Configuration > SSL > CA Certificates > Cached Intermediate Certificates**.
2. Select **Turn Caching Off** and click **OK** to confirm your decision.
3. Click **Apply**.

View Cached Intermediate Certificates

You can view information about the CA certificates, which conform to, at a minimum, the standards established within the PKI ITU-T X.509 standard.

To view the details of a specific cached intermediate certificate:

1. Select **Configuration > SSL > CA Certificates > Cached Intermediate Certificates**.
2. Select the cached intermediate certificate that you wish to see the details for and click **View**. Three certificate information tabs are available for analysis:
 - **General**—Includes top level information about a digital certificate, including the DN (distinguished name) identifying the owner and issuer, the dates when the certificate is valid, and the public key fingerprints using MD5 and SHA-1 cryptographic hash functions.
 - **Details**—Includes certificate field information as defined in the ITU-T X.509 public key certificate standard.
 - **PEM (Privacy-enhanced mail)**—Displays the certificate contents in a Base64 encoded format. You can copy the contents of the certificate to your clipboard by clicking on **Copy To Clipboard**.
3. Click **Close** when you have finished examining the contents.

Clear Cached Intermediate Certificates

Clearing the CA certificate cache removes all stored intermediate CA certificates.

To clear the cached intermediate certificates:

1. Select **Configuration > SSL > CA Certificates > Cached Intermediate Certificates**.
2. Select **Clear Cache** and click **OK** to confirm your decision.
3. Click **Apply**.

Note: The ProxySG appliance retains the list of cached intermediate CA certificates even after the appliance is shutdown and restarted. The only way to delete the cache is to manually clear or turn off certificate caching.

Section F: Checking Certificate Revocation Status in Real Time (OCSP)

This section describes how to use the Blue Coat ProxySG for performing real time certificate revocation checks using the Online Certificate Status Protocol (OCSP).

See Also

["About OCSP" on page 1151](#)

["How Blue Coat ProxySG Uses OCSP" on page 1151](#)

["OCSP CPL Policy Configuration" on page 1161](#)

["OCSP Listed Exceptions" on page 1161](#)

["OCSP Access Log Fields" on page 1161](#)

About OCSP

OCSP (RFC 2560) allows you to obtain the revocation status of an X.509 digital certificate. OCSP provides the same revocation functionality as the local Certificate Revocation List (CRL) configured on the ProxySG.

Managing large CRLs poses scalability challenges. This is due to high memory consumption on the ProxySG associated with storing revocation lists. OCSP overcomes these limitations by checking certificate status in real time using off-box OCSP responders.

How Blue Coat ProxySG Uses OCSP

The ProxySG can act as an OCSP client and issue OCSP queries to remote OCSP responders located on the intranet or the Internet. OCSP configuration and administration is usually performed by the administrator who manages the web access policy for an organization.

The ProxySG supports OCSP based revocation checks for:

- ❑ SSL proxy
- ❑ HTTPS reverse proxy
- ❑ SSL health checks
- ❑ secure image downloads
- ❑ secure URL database downloads
- ❑ secure heartbeats

OCSP-based revocation checks are performed on client or server certificates by the above applications where suitable. In this section, these client or server certificates are referred to as *subject certificates*. The ProxySG acts as an OCSP client and sends OCSP queries to an OCSP responder for the given certificate. An OCSP responder is a server for OCSP request processing and response building functions.

The OCSP responder sends status of the certificate back to the ProxySG (OCSP client). Status can be good, revoked or unknown. *Good* means that the certificate is not revoked and valid at the time of the query. *Revoked* means that the certificate

has been revoked either permanently or temporarily. *Unknown* means that the responder does not know about the revocation status of the certificate being requested.

The ProxySG can also cache OCSP responses and has the ability to respect, override or ignore the timestamps related to cacheability in the OCSP response.


If the certificate status is valid, the end user (in cases of SSL proxy or HTTPS reverse proxy) can access the secure website. If the status is revoked, an error is flagged and the end user is denied access to the secure website. If status is unknown, the ProxySG has the ability to treat it as an error or ignore it based on the administrator's discretion.

Basic OCSP Setup Scenarios

This section describes three general OCSP setup scenarios which are based on the relationship between the subject certificate (the certificate whose revocation status is queried, for example, client or server certificate) and the responder certificate (the certificate that signed the OCSP response).


In the following scenario illustrations, the subject certificate chain is comprised of certificates shown on the left-hand side. The responder certificate chain is comprised of certificates shown on the right-hand side.

Scenario A




The OCSP response is signed by a root CA that also issued the subject certificate.

Scenario B



The OCSP response is signed by a *delegated* certificate and both the responder certificate and the subject certificate are issued by the same root CA. The root CA in this scenario delegates the job of the signing OCSP responses to the OCSP responder by adding the OCSP signing purpose to the extendedKeyUsage extension of the responder's certificate (See section 4.2.2.2 of RFC 2560). This denotes that the certificate has been delegated for the purpose of signing OCSP responses by the root CA certificate.

Scenario C



The OCSP response is signed by a certificate having no common issuer with the subject certificate. Thus, the root CA certificates signing the subject certificate and OCSP response are different. This only works if the responder certificate's root CA is trusted by the administrator for the OCSP signing. The administrator can denote this trust by adding the `OCSP Signing` trust setting in the `Trusted Uses` section of the root CA. OpenSSL provides a command line tool to add this trust setting to a traditional root CA certificate.

Here is an example of how to create a root CA trusted for OCSP signing from an existing root:

```
openssl x509 -in <root CA file> -addtrust OCSPSigning -out  
<trusted root CA>
```

A trusted certificate is an ordinary certificate that has several additional pieces of information attached to it. Information can include the permitted and prohibited uses of the certificate and an alias. Trust settings are a non-standard way to override the purposes present in the `keyUsage` or `extendedKeyUsage` extensions of a certificate.

By default, a trusted certificate must be stored locally and must be a root CA. Trust settings currently are only used with a root CA. They allow finer control over the purposes for which the root CA can be used for. For example, a CA may be trusted for an SSL client but not SSL server use. Other trust values that are supported by OpenSSL include:

- ❑ `clientAuth` (SSL client use)
- ❑ `serverAuth` (SSL server use)
- ❑ `emailProtection` (S/MIME email)

Notes

- ❑ The keyword `TRUSTED` is denoted in the certificate header and footer:

```
-----BEGIN TRUSTED CERTIFICATE-----  
-----END TRUSTED CERTIFICATE-----
```
- ❑ The **Ignore OCSP signing purpose check** option (see [Step 5](#) on page 1159 in "[Creating and Configuring an OCSP Responder](#)") lists the errors that are related to the OCSP signing delegation. This applies to Scenarios B and C only.

Symantec Reverse Proxy and SSL Proxy Scenarios

Reverse Proxy Scenario

The following diagram shows how the ProxySG uses OCSP in a typical HTTPS reverse proxy scenario.




Figure 61–1 Reverse Proxy Scenario


SSL Proxy Scenario

In a common SSL proxy scenario, the ProxySG reads in the server certificate and sends an OCSP request to the responder to validate the certificate. Then based on the certificate status in the OCSP response the ProxySG denies or allows user access to content on the origin content server.


Creating and Configuring an OCSP Responder

To enable an OCSP revocation check, configure an OCSP responder profile:

1. Select the **Configuration > SSL > OCSP** tab.



2. Click **New** to create a new OCSP responder. The **Create OCSP responder** dialog displays.



3. Configure the OCSP responder options:
 - a. **Name**—Give the responder a meaningful name. If you are editing an existing responder, this field is grayed out.
 - b. **URL**—Indicates the location of the OCSP responder. The ProxySG needs this URL to locate the responder. This location can be obtained from the certificate's Authority Information Access (AIA) extension or from a user-defined configuration. The default is to use the URL from the certificate.
 - **Use URL from certificate**—Select this option if you want the ProxySG to look up the OCSP server location from the subject certificate's AIA extension.
 - **Use URL:**—Select this option if the location of the designated OCSP responder is known to you. Enter a specific responder HTTP or HTTPS URL.

- c. **Issuer CCL**—This option is used to decide which responder is contacted for a given client or server certificate. Typically each certificate issuer uses a designated OCSP responder for all the certificates it issues. The issuer CCL attribute allows the administrator to specify the certificate authorities (issuers) for which the responder in question is the designated responder. This means that when a certificate is signed by one of the CAs in this CCL, the OCSP query for that certificate will be sent to this responder.

In the section "[Basic OCSP Setup Scenarios](#)" on page 1152, the entire certificate chain shown on the left-hand side (including the root CA certificate) in each figure (except for the certificate appearing lowest in the chain) must be part of the issuer CCL. The left-hand side certificate chain represents the subject certificate chain, that is, certificates on which an OCSP query is done. OCSP revocation check happens for each certificate in the chain, including the root CA. If any CA in that chain is absent from the issuer CCL this responder will not be used to query the missing CA's OCSP status.


From the drop-down list, select a CA Certificate List (CCL) that contains the CA certificate names for which this is the designated responder. Each CA may only appear in one responder's Issuer CCL. The default is None. Thus, for a given certificate, this CCL is used to determine which responder to use when doing an OCSP check.

- d. **Response CCL**—This attribute is used during verification of OCSP responses. In the section "[Basic OCSP Setup Scenarios](#)" on page 1152, the entire certificate chain shown on the right-hand side (including the root CA certificate) in each figure (except for the certificate appearing lowest in the chain) must be part of this CCL. The right-hand side certificate chain represents all certificates in the signing hierarchy of the OCSP responder certificate. If any CA in that chain is absent from this CCL, then response verification fails and an untrusted-responder error is stored in the ProxySG event log.

From the drop-down list, select the CCL list you want to use. The default value is **browser-trusted**.


For Scenarios A and B, this CCL must contain the Root CA as depicted in the respective figures. For Scenario C, the CCL must contain at least the Root CA. The root CA must be imported on the ProxySG using the trusted certificate format (with OCSPSigning trust enabled). If OCSP responder does not chain all intermediate CAs, then this CCL must also include all those intermediate CAs, otherwise an untrusted-responder error is stored in the event log.

- e. **Device Profile**—This attribute is used when the responder URL is an HTTPS URL. From the drop-down list, select the device profile you want to use when connecting to the OCSP server via SSL. All existing profiles on the ProxySG appear. The device profile is a unique set of SSL cipher-suites, protocols, and keyrings. When the responder URL is HTTPS the ProxySG makes the HTTPS connection with this responder using its device profile. If the URL is HTTP the device profile is not used. The default value for the device profile attribute is **default**.
- f. **Response Cache TTL**—This option indicates how many days an OCSP response is cached on the ProxySG. The default is to use TTL from OCSP response.
 - **Use the TTL from OCSP response**—Select this option to use the value of nextUpdate timestamp (see section 2.2 of RFC 2560) in the OCSP response. If this timestamp is not set or is in the past, the OCSP response is not cached on the ProxySG. The ProxySG permits a clock skew of up to five minutes with the responder's clock when validating the nextUpdate timestamp.
 - **Use the TTL**—Enter the length of time (in days) you want the OCSP response to be cached regardless of nextUpdate timestamp in the OCSP response. If TTL is set to 0, the response is not cached.
- g. **Enable forwarding**—This option specifies that OCSP requests are to be sent through a forwarding host, if configured. The default is to have forwarding enabled. Based on whether the responder URL is HTTP or HTTPS the usual forwarding rules apply.



4. Configure the extensions options:
 - a. **Enable nonce**—To avoid replay attacks, click **Enable nonce**. A *nonce* is a random sequence of 20 bytes placed in an OCSP response. The default is to disable the use of a nonce.
 - b. **Request signing keyring**— This keyring is used when an OCSP request is required to be signed. In this case, the ProxySG includes the certificate chain (minus the root CA) that is associated with this keyring to help the OCSP responder verify the signature.

When a valid keyring is selected then OCSP request signing is enabled. When **None** is selected no request signing occurs.



5. Configure the following **Ignore Settings**:

- **Ignore request failures**—This setting ignores various connection errors. By default, connection errors are not ignored. The following failures are ignored by this setting:
 - The responder's URL is set to `from-certificate` and the URL in the certificate's AIA extension is neither HTTP or HTTPS, or is not a valid URL.
 - The TCP layer fails to connect with the responder.
 - The responder URL is HTTPS and the initial SSL connection fails with the responder.
 - The TCP connection times out while reading the response from the responder.
 - The TCP connection fails for any reason not already listed.
 - The responder URL is HTTPS and a hostname mismatch error occurs on the responder's certificate.
 - The responder URL is HTTPS and an error occurs while analyzing the response. Any other error not caught is covered by the following ignore settings.
 - The OCSP responder returns an error message that is described in section 2.3 of RFC 2560. For instance, when an OCSP query is sent to a responder that is not authorized to return an OCSP status for that certificate, the responder returns an `unauthorized` error, that appears as `Responder error (unauthorized)` in event-log of the ProxySG. Enabling this setting causes this error to be ignored as well as other errors described in the RFC.
 - The OCSP responder returns a response that is not a basic OCSP response (see section 4.2.1 of RFC 2560).
- **Ignore expired responder certificate**—This setting ignores invalid dates in the responder certificate. By default, invalid responder certificate dates cause the subject certificate verification to fail.
- **Ignore untrusted responder certificate**—This setting ignores the response validation error that occurs when the responder's certificate cannot be trusted. By default, any untrusted certificate failure is an error and causes subject certificate verification to fail.

- **Ignore OCSP signing purpose check**—This setting ignores errors which are related to the OCSP signing delegation and applies only to Scenarios B and C. (See "[Basic OCSP Setup Scenarios](#)" on page 1152.) The errors might occur in one of two ways:
 - Scenario B—The response signer certificate is not delegated for the OCSP signing. The event log records this error as `missing ocspsigning usage`.
 - Scenario C—The root CA does not have the trust setting enabled for the OCSP Signing. The event log records this error as `root ca not trusted`.

Either of these errors may be ignored by enabling this setting.


- **Ignore unknown revocation status**—Select this setting to ignore `unknown` revocation status as an error. By default, `unknown` status is an error and causes subject certificate verification to fail.

6. Click **OK**.
7. Click **Apply**.

Setting the Default Responder

To set the default responder OCSP responder profile:

1. Select the **Configuration > SSL > OCSP** tab.



2. From the **Default Responder** drop-down list, select the responder you want to be designated as the default responder. If a responder has not been previously created then **<None>** is the only option.

If the subject certificate is not associated with any responder (using Issuer CCL option) then the OCSP request for this certificate is sent to the default responder.

Important: If the default responder has a URL that is set to `from certificate` (see Step 3b in "[Creating and Configuring an OCSP Responder](#)" on page 1155), then all ProxySG components which are capable of performing OCSP checks generate OCSP requests to responders that may be anywhere on the Internet depending on where the certificate's AIA extension URL is pointing. Use a default responder that has its URL set to `from certificate` with caution.

3. Click **Apply**.

OCSP CPL Policy Configuration

The following policy property is extended for revocation check under the SSL layer:

```
<ssl>
    server.certificate.validate.check_revocation(auto|local|ocsp|no)
</ssl>
    client.certificate.validate.check_revocation(auto|local|ocsp|no)
```

For detailed information about CPL policy configuration and revocation check, refer to the *Content Policy Language Guide*.

OCSP Listed Exceptions

When a certificate state is revoked, the following predefined exceptions are sent depending on which certificate is revoked:

- ☐ `ssl_client_cert_revoked`
- ☐ `ssl_server_cert_revoked`

When a certificate status is unknown and the responder is configured to not ignore it, the following predefined exceptions are sent depending on which certificate is revoked:

- ☐ `ssl_client_cert_unknown`
- ☐ `ssl_server_cert_unknown`

For detailed information about defining Exceptions, refer to the *Visual Policy Manager Reference*, Chapter 4, Table 4-1.

OCSP Access Log Fields

Note: See [Chapter 29: "Creating Custom Access Log Formats"](#) on page 649 for detailed information about creating and editing log formats.

The following table lists and describes the OCSF access log fields:

Table 61–1 Access Log Substitutions

ELFF	Description
x-rs-ocsp-error	An error was observed during the OCSF check for a server certificate.
x-cs-ocsp-error	An error was observed during the OCSF check for a client certificate.

The OCSF access log field descriptions are:

Table 61–2 Access Log Field Descriptions

Access Log Field	Description
unsupported-responder-url	An error occurs if: - The responder's URL is set to <code>from-certificate</code> and the URL in the target certificate's AIA field is neither HTTP or HTTPS. - Or, the URL is not a valid.
connection-failure	An error occurs during the TCP connection with the responder.
ssl-handshake-error	An error occurs over the HTTPS transport during the initial SSL handshake with the responder.
request-timeout	An error occurs if the TCP times out while reading the response from responder.
connection-dropped	An error occurs when any other TCP failure happens which is not encountered in the errors described in this table.
ssl-cert-hostname-mismatch	An error occurs during the HTTPS transport when there is a hostname mismatch on the responder front-end certificate.
invalid-response	An error occurs during the parsing of an OCSF response. For example, during an HTTP parsing error.

Table 61–2 Access Log Field Descriptions

Access Log Field	Description
ocsp-signing-purpose-error	An error occurs during the OCSP response verification in the following cases (Refer to RFC 2560, section 4.2.2.2): • The response-signer's certificate's extendedKeyUsage does not have an OCSPsigning value making the signer unauthorized. • Or, the root certificate of the response-signer has the same missing extension value as above.
untrusted-responder-cert	An error occurs during response verification when the response signer's certificate is not trusted by ProxySG
expired-responder-cert	An error occurs during response verification when the response signer's certificate has invalid dates.
internal-error	An error occurs when an error happens that is not described in this table.

Chapter 62: Managing SSL Traffic

This section describes how to configure the SSL client and devices profiles, which are required for secure connections. These profiles are configured to group together the collection of settings required for an SSL connection. The profiles themselves include:

- ❑ Keyrings
- ❑ CA certificates
- ❑ CA Certificate List (CCL)
- ❑ Cipher Suite

CA certificates, keyrings, CCLs and cipher suites must be configured individually before being added to an SSL client profile or an SSL device profile. Except for cipher suites, discussed in ["Changing the Cipher Suite of the SSL Client"](#) on page 1167, these settings are discussed in greater detail in [Chapter 61: "Managing X.509 Certificates"](#) on page 1113.

This section discusses the following topics:

- ❑ [Section A: "SSL Client Profiles"](#) on page 1166.
- ❑ [Section B: "SSL Device Profiles"](#) on page 1170.
- ❑ [Section C: "Notes and Troubleshooting"](#) on page 1171.

Section A: SSL Client Profiles

This section discusses SSL Client profiles.

About the SSL Client Profile

The SSL client profile contains the settings needed to make an SSL connection; this profile can be used by any HTTP or HTTPS proxy service that needs to make an upstream SSL connection.

Note: The SSL proxy, also known as the SSL forward proxy, uses parameters taken from the SSL connection made by the client when originating SSL connections to the server. As a result, settings in the default SSL client profile are not applied to these connections.

To modify any parameters for SSL connections, change the corresponding SSL device-profile. You will need to modify the SSL client profile settings in the reverse proxy scenario only. This is because the reverse proxy uses the SSL client, instead of the SSL device profile, when connecting to the upstream OCS using HTTPS.

Default settings for the SSL client are:

- ☐ Keyring: None
- ☐ SSL Versions: TLSv1, TLSv1.1, TLSv1.2
- ☐ CCL: browser-trusted
- ☐ Cipher suite: All

Editing an SSL Client

The SSL client settings are global, affecting all services that use it. Unless required by your environment, you do not need to change any settings.


To change the protocol, the cipher suite, the keyring or the CCL associated with the SSL client, continue with "[Associating a Keyring, Protocol, and CCL with the SSL Client](#)" on page 1166 or "[Changing the Cipher Suite of the SSL Client](#)" on page 1167.

Associating a Keyring, Protocol, and CCL with the SSL Client

The SSL client, called default, already exists on the ProxySG.

To edit the SSL client:

1. Select **Configuration > SSL > SSL Client**.



2. Complete the following steps:
 - a. If the server in question requires a client certificate, select the keyring used to negotiate with origin content servers through an encrypted connection. Only keyrings with certificates can be associated with the SSL client, displayed in the **Keyring** drop-down list. By default, no keyring is selected.
 - b. (Optional) Change the **SSL Versions** default from **TLSv1.2, TLSv1.1, TLSv1** to any other combination of protocols listed in the list.
 - c. Select the CCL that the ProxySG uses to determine which CA certificates are trusted during server certificate validation. The CCL can be any already created certificate list. By default, the **browser-trusted** CCL is used.
3. Click **Apply**.

Changing the Cipher Suite of the SSL Client

The cipher suite sets the encryption method for the ProxySG. Changing the cipher suite can be done only through the CLI.

To change the cipher suite of the SSL client:

The default is to use all ciphers.

Note: Director uses non-interactive commands (commands that do not send options to the screen and wait for user input) to create the cipher suite used in Director overlays and profiles. For more information on Director, refer to the *Blue Coat Director Configuration and Management Guide*.

To change the cipher suite:

1. Select the ciphers you want to use at the prompt.

```
SGOS#(config) ssl
SGOS#(config ssl) edit ssl-client default
SGOS#(config ssl ssl-client default) cipher-suite
```

SSL-Client Name	Keyring Name	Protocol
-----	-----	-----
default	default	SSLv2v3TLSv1

Cipher#	Use	Description	Strength
-----	---	-----	-----
1	yes	AES128-SHA256	High
2	yes	AES256-SHA256	High
3	yes	AES128-SHA	Medium
4	yes	AES256-SHA	High
5	yes	DHE-RSA-AES128-SHA	High
6	yes	DHE-RSA-AES256-SHA	High
7	yes	DES-CBC3-SHA	High
8	yes	RC4-SHA	Medium
9	yes	RC4-MD5	Medium
10	yes	DES-CBC-SHA	Low
11	yes	EXP-DES-CBC-SHA	Export
12	yes	vEXP-RC4-MDS	Export
13	yes	EXP-RC2-CBC-MD5	Export

```
Select cipher numbers to use, separated by commas: 1,3,4
ok
```

2. (Optional) View the results.

```
SGOS#(config ssl ssl-client default) view
SSL-Client Name: default
Keyring Name: <None>
CCL: browser-trusted
Protocol: tlsv1 tlsv1.1 tlsv1.2
Cipher suite: 123128-sha256 aes256-sha256 aes128-sha aes256-sha dhe-
rsa-aes128-sha dhe-rsa-aes256-sha des-cbc3-sha rc4-sha rc4-md5 des-
cbc-sha exp-des-cbc-sha exp-rc4-med5 exp-rc2-cbc-md5
```

To change the cipher suite non-interactively:

Enter the following commands:

```
SGOS#(config) ssl
SGOS#(config ssl) edit ssl-client default
SGOS#(config ssl ssl-client default) cipher-suite cipher
```

where *cipher* is any of those listed above

Notes:

- ❑ If you do not specify any attributes, the cipher suite cannot be used.
- ❑ Multiple ciphers can be specified on the command line, separated by blank spaces.

Example

```
SGOS#(config ssl ssl-client default) cipher-suite rc4-sha
ok
SGOS#(config ssl ssl-client default) view
SSL-Client: default
Keyring: <None>
CCL: browser-trusted
Protocol: SSLv2v3TLSv1
Cipher suite: rc4-sha
```

Section B: SSL Device Profiles

This section discusses SSL Device profiles.

About SSL Device Profiles

An SSL device profile contains the settings needed to make an SSL connection to a remote system; this profile is used when the ProxySG is an SSL endpoint for non-proxy traffic, such as secure ADN connections, LDAP client, BCAA client, and WebPulse. The ProxySG is pre-configured with three SSL device profiles, each with a specific purpose. You can create other profiles for other purposes or edit the default profile to suit the environment.

To modify any parameters for SSL connections, change the corresponding SSL device-profile except in the case of the reverse proxy scenario. This is because the reverse proxy uses the SSL client, instead of the SSL device profile, when connecting to the upstream OCS using HTTPS.

Note: Non-proxy traffic uses an SSL device profile. Proxy traffic uses the SSL client profile. For proxy traffic, see [Section A: "SSL Client Profiles"](#) on page 1166.

The already-created SSL device profiles and their purposes are:

- ❑ **bluecoat-appliance-certificate:** This profile, which cannot be edited or deleted, is used for device-to-device authentication, allowing Blue Coat devices on a network to identify other Blue Coat devices that can be trusted. You can select this device profile when setting up device authentication, or you can create a new device profile as described in ["Creating an SSL Device Profile for Device Authentication"](#) on page 1296.
- ❑ **passive-attack-detection-only:** This profile, which cannot be edited or deleted, optionally can be used in place of the **bluecoat-appliance-certificate** profile. The **passive-attack-detection-only** profile uses a self-signed certificate and disables the verify-peer option, so that no authentication is done on the endpoints of the connection. The traffic is encrypted, but is vulnerable to active attacks.
- ❑ **default:** This profile can be edited but not deleted. Only secured non-proxy traffic uses this profile.

Some non-proxy traffic, such as ADN, has no default profile; you must choose a profile before enabling security for the traffic.

Editing or Creating an SSL Device Profile

You can edit the existing default SSL device profile for the environment and also create additional SSL device profiles with different settings. For example, if you require a different cipher setting from what the default profile uses, create a profile with the different cipher suite.

For instructions, see ["Creating an SSL Device Profile for Device Authentication"](#) on page 1296.

Section C: Notes and Troubleshooting

The following topics apply to both the SSL Client and the SSL device profiles.

Troubleshooting Server Certificate Verification

The three most common causes of server certificate verification failure are:

- ❑ The absence of a suitable CA certificate on the ProxySG. Be sure that the ProxySG is configured with the relevant CA certificates to avoid unwanted verification failures.
- ❑ The certificate is being used before its valid-from date or used after its valid-to date. This generally happens when a clock mismatch occurs between the certificate and the machine using the certificate. It is also possible that the clock on one of the machines is wrong.
- ❑ The common name in the certificate might not match the hostname in the URL.

Server certification validation can also be controlled through policy:

- ❑ CPL: Use the `server.certificate.validate( )` property in the Forwarding layer.
- ❑ VPM: Use the **Set Server Certificate Validation** action in the SSL Access Layer.

Setting the SSL Negotiation Timeout

The SSL negotiation timeout value dictates the time a ProxySG waits for a new SSL handshake to complete.

You can change the default SSL negotiation timeout value if the default, 300 seconds, is not sufficient for the environment. This value can only be changed through the CLI; it cannot be set from the Management Console.

To change the timeout period, enter the following commands from the command prompt:

```
SGOS#(config) ssl
SGOS#(config ssl) view ssl-nego-timeout
300
SGOS#(config ssl) ssl-nego-timeout seconds
```


Chapter 63: Windows Single Sign-on Authentication

This section describes how to configure the Windows Single Sign-on (SSO) realm, which is an authentication mechanism available on Windows networks. It includes the following topics:

- ❑ ["How Windows SSO Realms Work"](#) on page 1175
- ❑ ["Creating a Windows SSO Realm"](#) on page 1178
- ❑ ["Configuring Windows SSO Agents"](#) on page 1178
- ❑ ["Configuring Windows SSO Authorization"](#) on page 1180
- ❑ ["Defining Windows SSO Realm General Properties"](#) on page 1182
- ❑ ["Modifying the sso.ini File for Windows SSO Realms"](#) on page 1183
- ❑ ["Creating the CPL"](#) on page 1185
- ❑ ["Notes"](#) on page 1186

How Windows SSO Realms Work

In a Windows SSO realm, the client is never challenged for authentication. Instead, the BCAA agent collects information about the current logged on user from the domain controller and/or by querying the client machine. Then the IP address of an incoming client request is mapped to a user identity in the domain. If authorization information is also needed, then another realm (LDAP or local) must be created. For more information, see ["How Windows SSO Authorization Works"](#) on page 1177.

Windows SSO realms are compatible with administrative authentication configurations, but not recommended because they do not challenge the user to authenticate. Windows SSO relies on the LDAP server to identify the user requesting access based on their client IP address.

Note: The Windows SSO realm works reliably only in environments where one IP address maps to one user. If an IP address cannot be mapped to a single user, authentication fails. Those with NAT systems, which uses one set of IP addresses for intranet traffic and a different set for Internet traffic, should use a different realm for authentication

To authenticate a user, the Windows SSO realm uses two methods, either separately or together:

- ❑ **Domain Controller Querying:** The domain controller is queried to identify which users are connecting to, or authenticating with, the domain controller. This can be used to infer the identity of the user at a particular workstation.

- ❑ Client Querying: The client workstation is queried to determine who the client workstation thinks is logged in.
- ❑ When Domain Controller Querying and Client Querying are both used, the Domain Controller Query result is used if it exists and is still within the valid time-to-live as configured in the `sso.ini` file. If the Domain Controller Query result is older than the configured time-to-live, the client workstation is queried.

Note: The BCAAA 6.0 installer automatically enables Domain Controller Query (DCQ) in the `sso.ini` file when the user indicates that they will use Windows SSO. To enable DCQ in earlier BCAAA releases, you must modify the `sso.ini` file (located in the same directory as the BCAAA service). For information on modifying this file, see ["Modifying the sso.ini File for Windows SSO Realms"](#) on page 1183.

For the most complete solution, an IWA realm could be configured at the same time as the Windows SSO realm and both realms added to a realm sequence. Then, if the Windows SSO realm failed to authenticate the user, the IWA realm could be used. For information on using a sequence realm, see [Chapter 60: "Sequence Realm Authentication"](#) on page 1107.

Administrative authentication with Windows SSO is insecure, as the user is not challenged to authenticate when accessing the ProxySG appliance management console. For this reason, Symantec recommends Local or Certificate realms, or IWA with BCAAA secured over TLS for administrative authentication.

How Windows SSO Works with BCAAA

The server side of the authentication exchange is handled by the Blue Coat Authentication and Authorization Agent (BCAAA). Windows SSO uses a single BCAAA process for all realms and proxies that use SSO.

BCAAA must be installed on a domain controller or member server. By default, the BCAAA service authenticates users in all domains trusted by the computer on which it is running. When using Domain Controller Querying, the BCAAA service can be configured to only query certain domain controllers in those trusted domains.

Note: Refer to the *BCAAA Service Requirements* document for up-to-date information on BCAAA compatibility. The *BCAAA Service Requirements* document is posted at MySymantec.

By default the BCAAA service is installed to run as LocalSystem. For a Windows SSO realm to have correct permissions to query domain controllers and clients, the user who BCAAA runs under must be an authenticated user of the domain.

When the Windows SSO realm is configured to do Client Querying, the user that BCAAA runs under must be an authenticated user of the domain. For failover purposes, a second BCAAA can be installed and configured to act as an alternate BCAAA in the Windows SSO realm. The alternate BCAAA service is used in the event of a failure with the primary BCAAA service configured in the realm.

BCAAA Synchronization

Optionally, when using Domain Controller Querying, you can configure a BCAAA service to use another BCAAA service as a synchronization server. Whenever a BCAAA service restarts, it contacts its synchronization server and updates the logon state. Two given BCAAA services can use each other as their synchronization server. Thus, each BCAAA service can act as a synchronization server to provide logon state to other BCAAA services, as well as acting as a synchronization client to update its logon state from another BCAAA service.

Each BCAAA service has a synchronization priority that determines synchronization behavior. If the client BCAAA has the same or higher priority than the server BCAAA, synchronization is done once at restart to update the client state. Once synchronization is complete the client BCAAA drops the synchronization connection and begins querying the domain controllers.

However, if the server BCAAA has higher priority, then the client BCAAA keeps the synchronization link open and continuously updates its logon state from the higher priority BCAAA. The client BCAAA does not query the domain controllers itself unless the synchronization link fails.

This makes it possible to manage the query load on the domain controllers. If there is no issue with load, then the default configuration (without synchronization), with all BCAAA agents querying the domain controllers is acceptable. However, if load on the domain controllers is an issue, synchronization can be used to minimize this load while still providing fail-over capabilities.

By default, all BCAAA agents have the same synchronization priority, meaning that they synchronize on startup and then do their own domain controller querying. To change the synchronization settings, see ["To configure the sso.ini file for synchronization:"](#) on page 1184.

Note: For information on configuring the BCAAA service as an authenticated user of the domain, refer to the *BCAAA Service Requirements* document posted at MySymantec.

How Windows SSO Authorization Works

The Windows SSO realm, in addition to allowing you to create and manipulate realm properties, such as the query type and the number of seconds that credential cache entries from this realm are valid, also contains the authorization username and the name of the realm that will do authorization for the Windows SSO realm. The authorization username is a string containing policy substitutions that describes how to construct the username for authorization lookups. This can either be an LDAP FQDN when the authorization realm is an LDAP realm, or a simple name when local realms are being used for authorization.

Note: Windows SSO realms never challenge for credentials. If the authorization username cannot be determined from the configured substitutions, authorization in the Windows SSO realm fails.


Windows SSO realms do not require an authorization realm. If no authorization realm is configured, the user is not considered a member of any group. The effect this has on the user depends on the authorization policy. If the policy does not make any decisions based on groups, you do not need to specify an authorization realm. Also, if your policy is such that it works as desired when all Windows SSO realm users are not in any group, you do not have to specify an authorization realm.

Creating a Windows SSO Realm

This section describes how to create an SSO realm.

To create a Windows SSO realm:

1. Select the **Configuration > Authentication > Windows SSO > Windows SSO Realms** tab.
2. Click **New**.



3. In the **Realm name** field, enter a realm name. The name can be 32 characters long and composed of alphanumeric characters and underscores. The name *must* start with a letter.
4. Click **OK**.
5. Click **Apply**.

Configuring Windows SSO Agents

You must configure the Windows realm so that it can find the Blue Coat Authentication and Authorization Agent (BCAAA).

1. Select **Configuration > Authentication > Windows SSO > Agents**.

2. Select the **Realm name** to edit from the drop-down list.
3. In the **Primary agent** area (**Host** field), enter the hostname or IP address where the BCAAA agent resides. Change the port from the default of **16101** if necessary.
4. (Optional) Enter an alternate agent host and agent name in the **Alternate agent** area (**Host** field). The primary and alternate BCAAA server must work together to support fail-over. If the primary BCAAA server fails, the alternate server should be able to provide the same mappings for the IP addresses.
5. (Optional) Configure SSL options:
 - a. Click **Enable SSL** to enable SSL between the ProxySG and BCAAA.
 - b. (Optional) Select the SSL device profile that this realm uses to make an SSL connection to a remote system. You can choose any device profile that displays in the drop-down list. For information on using device profiles, see "[Appliance Certificates and SSL Device Profiles](#)" on page 1290.
6. In the **Timeout Request** field, type the number of seconds the ProxySG allows for each request attempt before timing out. (The default request timeout is **60** seconds.)
7. In the **Query Type** field, select the method you want to use from the drop-down menu.

If all of the client computers can be queried directly, then the most accurate results can be provided by the **Query Clients** option.

By default the Windows SSO realm is configured for **Domain Controller Querying**.

Client Querying is blocked by the Windows XP SP2 firewall. This can be overridden through domain policy. If the firewall setting **Allow remote administration exception** or **Allow file and printer sharing exception** or **Define port exceptions (with port 445)** is enabled, then the query will work.

If an authentication mode without surrogate credentials is being used (Proxy or Origin authenticate mode), then the **Query Domain Controller and Client** and **Query Client** options can cause too much traffic when querying the clients, as each authentication request results in a request to the BCAAA service, which can result in a client workstation query depending on the client query time-to-live. If the client workstation querying traffic is a concern, the **Query Domain Controllers** option should be used instead.

8. Click **Apply**.
9. Verify the Windows SSO configuration as follows:
 - a. Click **Test Configuration**. The Test Configuration dialog displays.
 - b. Enter the **IP address** of a client system in your Active Directory and then click **OK**. The ProxySG appliance will use configuration you supplied to send an authentication request to BCAAA and return the results as follows:
 - If the ProxySG and the BCAAA server are configured properly, BCAAA will return the username associated with the IP address you provided.
 - If the test does not succeed, check that the settings on the **Agents** tab as well as the BCAAA settings are configured properly and then test the configuration again.

Configuring Windows SSO Authorization

After the Windows SSO realm is created, you can use the Windows SSO Authorization tab to configure authorization for the realm.

Note: Windows SSO realms do not require an authorization realm. If the policy does not make any decisions based on groups, you do not need to specify an authorization realm.

Prerequisite

You must have defined at least one Windows SSO realm (using the Windows SSO Realms tab) before attempting to set Windows SSO realm properties. If the message **Realms must be added in the Windows SSO Realms tab before editing this tab** is displayed in red at the bottom of this page, you do not currently have any Windows SSO realms defined.

1. Select the **Configuration > Authentication > Windows SSO > Authorization** tab.

The screenshot shows a configuration window with three tabs: 'Agents', 'Authorization', and 'Windows SSO General'. The 'Authorization' tab is active. It contains three fields: 'Realm name' (a dropdown menu showing 'WindowSSOTest'), 'Authorization realm name' (a dropdown menu showing 'LDAP_test'), and 'Authorization username' (a text field containing 'cn=\$(user.name),ou=partition,o=company'). Below these fields are two radio buttons: 'Use FQDN' (which is selected) and an unlabeled radio button.

2. Configure authorization options:
- a. From the **Realm name** drop-down list, select the Windows SSO realm for which you want to change realm properties.
 - b. (Optional) From the **Authorization realm name** drop-down list, select the previously-configured realm used to authorize users.

To construct usernames, remember that the authorization username attributes is a string that contains policy substitutions. When authorization is required for the transaction, the character string is processed by the policy substitution mechanism, using the current transaction as input. The resulting string becomes the user's authorization name for the current transaction.
 - c. By default, the LDAP FQDN is selected as the **Authorization user name**. Change this value if the user's authorization information resides in a different root DN. To use a different authorization name, de-select **Use FQDN** and enter a different name, for example:
`cn=$(user.name),ou=partition,o=company`

3. Click **Apply**.

Table 63–1 Common Substitutions Used in the Authorization username Field

ELFF Substitution	CPL Equivalent	Description
x-cs-auth-domain	<code>\$(user.domain)</code>	The Windows domain of the authenticated user.
cs-username	<code>\$(user.name)</code>	The relative username of the authenticated user.

Defining Windows SSO Realm General Properties

The **Windows SSO General** tab allows you to specify the display name, the refresh times, an inactivity timeout value, cookies, and a virtual URL.

Windows SSO realms default to the origin-ip authentication mode when either no authentication mode or the auto authentication mode is specified in policy. After a user has first successfully authenticated to the ProxySG, all subsequent requests from that same IP address for the length of the surrogate credential refresh time are authenticated as that user. If the first user is allowed or denied access, subsequent users during that same time coming from the same IP address are allowed or denied as that first user. This is true even if policy would have treated them differently if they were authenticated as themselves.

If multiple users often log in from the same IP address, it is recommended to use a shorter surrogate credential refresh timeout than the default or an authentication mode that uses cookie surrogate credentials.

Prerequisite

You must have defined at least one Windows SSO realm (using the Windows SSO Realms tab) before attempting to set Windows SSO general properties. If the message **Realms must be added in the Windows SSO Realms tab before editing this tab** displays in red at the bottom of this page, you do not currently have any Windows SSO realms defined.

To configure general settings:

1. Select the **Configuration > Authentication > Windows SSO > Windows SSO General** tab.

2. From the **Realm name** drop-down list, select the Windows SSO realm for which you want to change properties.
3. Configure refresh options:
 - a. Select the **Use the same refresh time for all** check box if you would like to use the same refresh time for all.

- b. Enter the number of seconds in the **Surrogate refresh time** field. The Surrogate Refresh Time allows you to set a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here. Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the ProxySG authenticates the transaction. After the refresh time expires, the ProxySG determines which user is using the current IP address, and update the surrogate credential to authenticate with that user.
 - c. Enter the number of seconds in the **Authorization refresh time** field. The Authorization Refresh Time allows you to manage how often the authorization data is verified with the authentication realm. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.
4. Enter the number of seconds in the **Inactivity timeout** field to specify the amount of time a session can be inactive before being logged out.
5. Configure cookie options:
 - a. Select the **Use persistent cookies** check box to use persistent browser cookies instead of session browser cookies.
 - b. Select the **Verify the IP address in the cookie** check box if you would like the cookies surrogate credentials to only be accepted for the IP address that the cookie was authenticated. Disabling this allows cookies to be accepted from other IP addresses.
6. You can specify a virtual URL. For more information on the virtual URL, see ["About Origin-Style Redirection"](#) on page 910.
7. Click **Apply**.

Modifying the sso.ini File for Windows SSO Realms

SGOS 6.5 and later: You do not need to modify the `sso.ini` file to enable DCQ. If you are using BCAAA 6.0, the installer automatically enables DCQ in the `sso.ini` file when the user indicates that they will use Windows SSO.

BCAAA 5.5.x and earlier: To enable the method of authentication querying you choose, you must modify the `sso.ini` file by adding domain controllers you want to query and user accounts you want to ignore.

The `sso.ini` file is located in the BCAAA installation directory.

If you are only using one method of querying, you only need configure the specific settings for that method. If you plan to use both methods to query, you must configure all the settings.

Note: The changes to the `sso.ini` file have no effect until the BCAAA service is restarted.

To configure the `sso.ini` file for Domain Controller Querying

1. Open the file in a text editor.
2. In the section `DCQSetup`, uncomment the line: `DCQEnabled=1`.
3. In the section `DCQSetup`, set the `ValidTTL` time to mark users as logged out after a defined number of seconds. This prevents stale mappings in the IP-to-user-table. For example, setting `ValidTTL` to 86400 requires users log into their workstations at least once per day in order to be considered logged in by the ProxySG.
4. In the section `DCQDomainControllers`, list the domain controllers you want to query or the IP address ranges of interest.

By default all domain controllers that are in the forest or are trusted are queried. In large organizations, domain controllers that are not of interest for the ProxySG installation might be queried. The `sso.ini` file can be used to list the domain controllers of interest or IP address ranges of interest.

5. In the section `SSOServiceUsers`, list the domain names of users who can access the domain controller on behalf of the service and mask the identity of the logged-on user.

Listing these users here forces the BCAAA service to ignore them for authentication purposes.

6. Save the `sso.ini` file.

To configure the `sso.ini` file for client querying:

Note: Before you use the Windows SSO realm, you must change the BCAAA service to run as a domain user, and, if using XP clients, update the domain policy to allow the client query to pass through the firewall.

For information on installing and configuring the BCAAA service, refer to the *BCAAA Service Requirements* document posted at MySymantec.

1. Open the file in a text editor.
2. Review the TTL times in the `ClientQuerySetup` section to be sure they are appropriate for your network environment.
3. Update the `SSOServiceUsers` section to ignore domain users used for services.
4. Save the `sso.ini` file.

To configure the `sso.ini` file for synchronization:

1. Open the file in a text editor.
2. Update the section `SSOSyncSetup` (the defaults are listed below). Note that explanations of each setting are provided in the `sso.ini` file.

- `ServerPriority=100`
 - `EnableSyncServer=1`
 - `SyncPortNumber=16102`
 - `UseSSL=0`
 - `VerifyCertificate=0`
 - `QueryDelta=10`
 - `RetrySyncTime=60`
3. Update the section `SSOSyncServer` with the IP address or hostname of the BCAAA service to use a synchronization server.
 4. In the section `SSOSyncClients`, list the IP addresses or hostnames of the BCAAA services that will use this BCAAA service as their synchronization service.
 5. Save the `sso.ini` file.

Creating the CPL

You can create CPL policies now that you have completed Windows SSO realm configuration. Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes.

The examples below assume the default policy condition is *allow*. On new systems, the default policy condition is *deny*.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file layers.

- ❑ Every Windows SSO-authenticated user is allowed access the ProxySG.

```
<Proxy>
  authenticate (WSSORealm)
```

- ❑ Group membership is the determining factor in granting access to the ProxySG.

```
<Proxy>
  authenticate (WSSORealm)
<Proxy>
  group="cn=proxyusers, ou=groups, o=myco" ALLOW
  deny
```

Using Single Sign-On Realms and Proxy Chains

Some Application Delivery Network (ADN) configurations mask the source IP address of the request. For example, if the path for a request is:

client workstation > branch proxy > data center proxy > gateway proxy

Policy running on the gateway might see the IP address of the data center proxy rather than the IP address of the client workstation.

Note: The source IP address is not masked if you use the **reflect client ip** attribute.

In this ADN configuration, policy needs to be configured so that Windows SSO, Novell SSO, and policy substitution realms can authenticate users correctly.

Use the `user.login.address` and `authenticate.credentials.address` policy gestures to override the IP address of the credentials used for authentication and match the IP address of the authenticated user.

Note: The `user.login.address` condition only works correctly if you use the `authenticate.credentials.address` property to set the address.

You can also use the `x-cs-user-login-address` substitution to log this event.

Examples

In the following example, the address to use for authenticating with **myrealm** is set to the address received from the `HTTP Client-IP` header.

```
<proxy>
  authenticate(myrealm) \
  authenticate.credentials.address($(request.header.Client-IP))
```

In the following example, the user is authenticated if logged in from the `1.2.3.0/24` subnet.

```
<proxy>
  user.login.address=1.2.3.0/24 allow
```

Notes

- ❑ The Windows SSO realm works reliably only in environments where one IP address maps to one user.
- ❑ This realm never uses a password.
- ❑ When doing domain controller querying, the Windows SSO realm can lose the logon if the NetBIOS computer name cannot be determined through a DNS query or a NetBIOS query. The DNS query can fail if the NetBIOS name is different than the DNS host name or if the computer is in a different DNS domain than the BCAAA computer and the BCAAA computer is not set up to impute different DNS domains.

The NetBIOS query can fail because the NetBIOS broadcast does not reach the target computer. This can happen if the computer is behind a firewall that is not forwarding NetBIOS requests or if the computer is on a subnet that is not considered to be local to the BCAAA server.

To prevent this issue, the BCAAA machine must be configured to be able to query the NetBIOS name of any computer of interest and get the correct IP address.

One workaround is to use a WINS server. This works like a DNS server but handles NetBIOS lookups.

Chapter 64: Using XML Realms

This section discusses XML realms, which are used to integrate SGOS with the authentication/authorization protocol. If you use an authentication or authorization protocol that is not natively supported by Blue Coat, you can use the XML realm.

Topics in this Section

This section includes information about the following topics:

- ❑ ["About XML Realms"](#)
- ❑ ["Before Creating an XML Realm" on page 1190](#)
- ❑ ["Creating an XML Realm" on page 1190](#)
- ❑ ["Configuring XML Servers" on page 1191](#)
- ❑ ["Configuring XML Options" on page 1192](#)
- ❑ ["Configuring XML Realm Authorization" on page 1193](#)
- ❑ ["Configuring XML General Realm Properties" on page 1194](#)
- ❑ ["Creating the CPL" on page 1197](#)
- ❑ ["Viewing Statistics" on page 1197](#)

About XML Realms

An XML realm uses XML messages to request authentication and authorization information from an HTTP XML service (the XML *responder* that runs on an external server). The XML realm (the XML *requestor*) supports both HTTP GET and HTTP POST methods to request an XML response. The XML messages are based on SOAP 1.2.

The XML responder service accepts XML requests from the ProxySG, communicates with an authentication or authorization server, and responds with the result. When the realm is used to authenticate users, it challenges for Basic credentials. The username and password are then sent to the XML responder to authenticate and authorize the user.

The XML realm can place the username and password in the HTTP headers of the request or in the body of the XML POST request. If the credentials are placed in the HTTP headers, the Web server must do the authentication and the XML service just handles authorization. If credentials are placed in the XML request body, the XML service handles both authentication and authorization.

XML messages must conform to the Blue Coat XML realm schema. This is an XML schema based on SOAP 1.2. The *SGOS 6.4 Release Notes* provide the current location for this file.

An authenticate request sends the credentials to the XML responder and optionally sends the groups and attributes referenced in policy. The XML responder can then authenticate the credentials. The response indicates if the user was successfully authenticated and also includes the user's groups and attributes if the XML responder is performing authorization.

An authorize request sends the authenticated username to the XML responder and optionally sends the groups and attributes referenced in policy. The response includes the user's groups and attributes.

XML realms are not compatible with administrative authentication to the ProxySG appliance management console.

Before Creating an XML Realm

The following list describes the tasks you must complete before creating an XML realm.

- ❑ Create an appropriate XML realm responder (one that is designed to talk to the Blue Coat XML realm protocol) and install it on an HTTP Web server. You can either create the responder yourself or have a third party create it, such as Blue Coat Professional Services.

To create the XML realm responder, see [Chapter 76: "XML Protocol"](#) on page 1429 for a description of the SOAP protocol. The XML responder must correctly conform to the protocol. The XML realm performance is dependent on the response time of the XML responder.


- ❑ Configure an HTTP server with appropriate authentication controls. The authentication service can either depend on the HTTP server to authenticate the credentials, or the service can authenticate them directly. If the HTTP server is used to authenticate the credentials, it must be set up to protect the service with HTTP Basic authentication.
- ❑ (Optional) Configure an alternate HTTP server for redundancy. The XML responder service must be installed on the alternate server.

Creating an XML Realm

To create an XML realm:

Before you create an XML realm, be sure to complete the tasks in ["Before Creating an XML Realm"](#) on page 1190.

1. In the Management Console, select the **Configuration > Authentication > XML > XML Realms** tab.
2. Click **New**.



3. In the **Realm Name** field, enter a realm name. The name can be 32 characters long, composed of alphanumeric characters and underscores. The name *must* start with a letter.
4. Click **OK** to close the dialog.
5. Click **Apply**.


Configuring XML Servers

You do not need to change these values if the default settings are acceptable.

After you have created an XML realm, go to the XML Servers page to change current default settings.

To configure XML server properties:

1. In the Management Console, select the **Configuration > Authentication > XML > XML Servers** tab.



2. From the **Realm Name** drop-down list, select the XML realm.
3. Configure the Responder options:
 - a. **Responder:** Select the XML responder service to configure—**Primary** or **Alternate**—from the drop-down list. **Primary** is the default. You can configure both responder services before clicking **Apply**.

- b. **Host:** This is the hostname or IP address of the HTTP server that has the XML service. You must specify a host. The **port** defaults to port 80.
- c. **Authenticate request path:** Enter the XML responder path for authentication requests.
- d. **Authorize request path:** Enter the XML responder path for authorization requests.
4. In the **timeout request** fields, enter the number of seconds for the system to wait for a request and the number of times for the system to retry a request. The default is not to retry a request.
5. Specify the **maximum number of connections to the responder**. The default is five connections.
6. (Optional) Select **One-time passwords** to integrate with a non-Blue Coat supported authentication service that uses one-time passwords.

Note: One-time passwords are passwords that become invalid as soon as they are used. The passwords are often generated by a token or program, although pre-printed lists are also used. Using one-time passwords ensures that the password cannot be used in a replay attack.

7. Click **Apply**.
8. Repeat the above steps for additional XML realms, up to a total of 40.

Configuring XML Options

You do not need to change these values if the default settings are acceptable.

With XML realms, you can place the username and password in the HTTP headers of the request or in the body of the XML POST request. If the credentials are placed in the HTTP headers, the Web server can do the authentication and the XML service can just handle authorization. If the credentials are placed in the XML request body, the XML service handles both authentication and authorization.

To configure XML options:

1. In the Management Console, select the **Configuration > Authentication > XML > XML Options** tab.

The screenshot shows the 'XML Options' configuration page. It includes the following fields and options:

- 2** Realm name: XMLTest
- 3** HTTP request method: GET
- User credential location: GET (selected), POST
- 4** Put user credentials for authentication in the HTTP header (selected radio button)
- Put user credentials for authentication in the request (unselected radio button)
- 5** Username parameter: username

2. From the **Realm name** drop-down list, select the XML realm.
3. Select the HTTP request method: **GET** or **POST**.
4. Select a user credential option:
 - If the HTTP server is integrated with the authentication system, the HTTP server can authenticate the credentials. Select the **Put user credentials for authentication in the HTTP header** radio button. However, if this does not provide enough flexibility, the XML responder can do authentication.
 - To have the XML responder service handle both authentication and authorization, select the **Put user credentials for authentication in the request** radio button.
5. Enter the username parameter in the **Username parameter** field. The default is **username**.
6. Click **Apply**.

Configuring XML Realm Authorization

You do not need to change these values if the default settings are acceptable.

After you have created the XML realm, you still must take into consideration how you will use authentication and authorization:

- ☐ Use an XML realm for both authorization and authentication.

The realm is used for authentication and uses itself for authorization.

- ☐ Use an XML realm for authentication another realm for authorization.

An XML realm can be used for authentication and use another realm for authorization. The authorization realm can be a Local realm, an LDAP realm or another XML realm.

- ☐ Use an XML realm as an authorization realm for another realm.

An XML realm can be used as an authorization realm for another realm that is doing authentication. The authentication realm can be a Certificate realm, a Policy Substitution realm, a Novell SSO realm, a Windows SSO realm or another XML realm.

In all cases, you must write policy to authenticate and authorize the users. For information on writing policy for an XML realm, see "[Creating the CPL](#)" on page 1197.

To configure XML authorization properties:

1. In the Management Console, select the **Configuration > Authentication > XML > Authorization** tab.

The screenshot shows the 'Authorization' tab in the SGOS Administration Console. The 'Realm name' is set to 'XMLTest'. The 'Authorization realm name' is set to '<None>' and the 'Self' checkbox is checked. The 'Authorization username' field is empty, and the 'Use full username' radio button is selected. The 'Default group' field is empty. The checkbox 'Send the groups and attributes of interest in the request' is checked. Arrows labeled 2a through 2e point to these specific fields.

2. From the **Realm name** drop-down list, select the XML realm.
 - a. **Authorization realm name:** If the XML realm is not doing authorization, select an authorization realm from the drop-down list. By default, the authorization realm name is **Self**.

Note: If **Self** is selected, the **Authorization realm name** drop-down list is unavailable. To make the **Authorization realm name** drop-down list active, clear the **Self** check box.

- b. **Authorization username:** The default is **Use full username**. Clear the **Use full username** option to use a different name or to use a policy substitution that generates a username.
 - c. **Default group:** The default is no groups are selected.
 - d. The **send the groups and attributes of interest in the request** option is selected by default. These are the groups and attributes that are used in policy.
3. Click **Apply**.

Configuring XML General Realm Properties

The XML General page allows you to indicate the realm's display name, the refresh times, an inactivity timeout value, cookies, and a virtual URL for this realm.

To configure general XML settings:

1. In the Management Console, select the **Configuration > Authentication > XML > XML General** tab.

XML Servers XML Options Authorization XML General

2 → Realm name: XMLTest

Display name: XMLTest

3 → Refresh Times: ☒ Use the same refresh time for all

Credential refresh time: 900 seconds

Surrogate refresh time: 900 seconds

Authorization refresh time: 900 seconds

4 → Inactivity timeout: 900 seconds

5 → Rejected credentials time: 1 seconds

6 → Cookies

☐ Use persistent cookies

☒ Verify the IP address in the cookie

7 → Virtual URL: www.cfauth.com/

8 → ☒ Challenge user after logout

2. Configure realm name information:
 - a. From the **Realm name** drop-down list, select the XML realm for which you want to change properties.
 - b. If needed, give the LDAP realm a display name. The default value for the display name is the realm name. The display name cannot be greater than 128 characters and it cannot be null.
3. Configure refresh options:
 - a. Select the **Use the same refresh time for all** check box if you would like to use the same refresh time for all.
 - b. Enter the number of seconds in the **Credential refresh time** field. The Credential Refresh Time is the amount of time basic credentials (username and password) are kept on the ProxySG. This feature allows the ProxySG to reduce the load on the authentication server and enables credential spoofing. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here. Before the refresh time expires, the ProxySG authenticates the user supplied credentials against the cached credentials. If the credentials received do not match the cached credentials, they are forwarded to the authentication server in case the user password changed. After the refresh time expires, the credentials are forwarded to the authentication server for verification.

- c. Enter the number of seconds in the **Surrogate refresh time** field. The Surrogate Refresh Time allows you to set a realm default for how often a user's surrogate credentials are refreshed. Surrogate credentials are credentials accepted in place of a user's actual credentials. The default setting is 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.

Before the refresh time expires, if a surrogate credential (IP address or cookie) is available and it matches the expected surrogate credential, the ProxySG authenticates the transaction. After the refresh time expires, the ProxySG verifies the user's credentials. Depending upon the authentication mode and the user-agent, this may result in challenging the end user for credentials.

The main goal of this feature is to verify that the user-agent still has the appropriate credentials.

- d. Enter the number of seconds in the **Authorization refresh time** field. The Authorization Refresh Time allows you to manage how often the authorization data is verified with the authentication realm. It has a default setting of 900 seconds (15 minutes). You can configure this in policy for better control over the resources as policy overrides any settings made here.
4. Enter the number of seconds in the **Inactivity timeout** field to specify the amount of time a session can be inactive before being logged out.
5. If you use Basic credentials and want to cache failed authentication attempts (to reduce the load on the authentication service), enter the number of seconds in the **Rejected Credentials time** field. This setting, enabled by default and set to one second, allows failed authentication attempts to be automatically rejected for up to 10 seconds. Any Basic credentials that match a failed result before its cache time expires are rejected without consulting the back-end authentication service. The original failed authentication result is returned for the new request.

All failed authentication attempts can be cached: Bad password, expired account, disabled account, old password, server down.

To disable caching for failed authentication attempts, set the **Rejected Credentials time** field to 0.

6. Select the **Use persistent cookies** check box to use persistent browser cookies instead of session browser cookies.
7. Select the **Verify the IP address in the cookie** check box if you would like the cookies surrogate credentials to only be accepted for the IP address that the cookie was authenticated. Disabling this allows cookies to be accepted from other IP addresses.
8. You can specify a virtual URL. For more information on the virtual URL, see ["About Origin-Style Redirection"](#) on page 910.
9. Click **Apply**.

Creating the CPL

This CPL example gives access to users who are authenticated in the XML realm called **eng_users** and who are in the group **waterloo**. You also can create policy for XML realms through VPM.

Note: For information on using policy, refer to the *Visual Policy Manager Reference* or *Content Policy Language Guide*.

```
<proxy>
  authenticate(eng_users)
</proxy>
realm=eng_users group=waterloo allow
```

Viewing Statistics

To view statistics for XML realms, select **Statistics > Authentication > User Logins**. Select an XML realm from the Realm drop-down list.

Chapter 65: Forms-Based Authentication

This chapter discusses forms-based authentication exceptions, which control what your users see during an authentication process. With forms-based authentication, you can set limits on the maximum request size to store and define the request object expiry time. You can also specify whether to verify the client's IP address against the original request and whether to allow redirects to the original request.

This chapter includes the following sections:

- ❑ ["About Authentication Forms"](#) on page 1199
- ❑ ["Configuring Forms-Based Authentication"](#) on page 1204
- ❑ ["Creating and Editing a Form"](#) on page 1204
- ❑ ["Setting Storage Options"](#) on page 1206
- ❑ ["Using CPL with Forms-Based Authentication"](#) on page 1207
- ❑ ["Troubleshooting Forms-Based Authentication"](#) on page 1208

About Authentication Forms

With forms-based authenticating, you can set limits on the maximum request size to store and define the request object expiry time. You can also specify whether to verify the client's IP address against the original request and whether to allow redirects to the original request.

You can:

- ❑ Specify the realm the user is to authenticate against.
- ❑ Specify that the credentials requested are for the ProxySG. This avoids confusion with other authentication challenges.
- ❑ Make the form comply with company standards and provide other information, such as a help link.

The authentication form (an HTML document) is served when the user makes a request and requires forms-based authentication. If the user successfully authenticates to the ProxySG, the appliance redirects the user back to the original request.

If the user does not successfully authenticate against the ProxySG and the error is user-correctable, the user is presented with the authentication form again.

Note: You can configure and install an authentication form and several properties through the Management Console and the CLI, but you must use policy to dictate the authentication form's use.

To create and put into use forms-based authentication, you must complete the following steps:

- ❑ Create a new form or edit one of the existing authentication form exceptions
- ❑ Set storage options
- ❑ Set policies

Three authentication forms are created initially:

- ❑ **authentication_form:** Enter Proxy Credentials for Realm `$(cs-realm)`. This is the standard authentication form that is used for authentication with the ProxySG.
- ❑ **new_pin_form:** Create New PIN for Realm `$(cs-realm)`. This form is used if you created a RADIUS realm using RSA SecurID tokens. This form prompts the user to enter a new PIN. The user must enter the PIN twice in order to verify that it was entered correctly.
- ❑ **query_form:** Query for Realm `$(cs-realm)`. This form is used if you created a RADIUS realm using RSA SecurID tokens. The form is used to display the series of yes/no questions asked by the SecurID new PIN process.

You can customize any of the three initial authentication form exceptions or you can create other authentication forms. (You can create as many authentication form exceptions as needed. The form must be a valid HTML document that contains valid form syntax.)

Each authentication form can contain the following:

- ❑ **Title** and sentence instructing the user to enter ProxySG credentials for the appropriate realm.
- ❑ **Domain:** Text input with maximum length of 64 characters. The name of the input must be `PROXY_SG_DOMAIN`, and you can specify a default value of `$(x-cs-auth-domain)` so that the user's domain is prepopulated on subsequent attempts (after a failure).

The input field is optional, used only if the authentication realm is an IWA realm. If it is used, the value is prepended to the username value with a backslash.

- ❑ **Username:** Text input with maximum length of 64 characters. The name of the input must be `PROXY_SG_USERNAME`, and you can specify a default value of `$(cs-username)` so the username is prepopulated on subsequent attempts (after a failure).
- ❑ **Password:** The password should be of type `PASSWORD` with a maximum length of 64 characters. The name of the input must be `PROXY_SG_PASSWORD`.
- ❑ **Request ID:** If the request contains a body, then the request is stored on the ProxySG until the user is successfully authenticated.

The request ID should be of type `HIDDEN`. The input name must be `PROXY_SG_REQUEST_ID`, and the value must be `$(x-cs-auth-request-id)`. The information to identify the stored request is saved in the request id variable.

- ❑ **Challenge State:** The challenge state should be of type `HIDDEN`. If a RADIUS realm is using a response/challenge, this field is used to cache identification information needed to correctly respond to the challenge.
The input name must be `PROXY_SG_PRIVATE_CHALLENGE_STATE`, and the value must be `$(x-auth-private-challenge-state)`.
- ❑ **Submit button.** The submit button is required to submit the form to the ProxySG.
- ❑ **Clear form button.** The clear button is optional and resets all form values to their original values.
- ❑ **Form action URI:** The value is the authentication virtual URL plus the query string containing the base64 encoded original URL `$(x-cs-auth-form-action-url)`.
- ❑ **Form METHOD of POST.** The form method must be POST. The ProxySG does not process forms submitted with GET.

The ProxySG only parses the following input fields during form submission:

- ❑ `PROXY_SG_USERNAME` (required)
- ❑ `PROXY_SG_PASSWORD` (required)
- ❑ `PROXY_SG_REQUEST_ID` (required)
- ❑ `PROXY_SG_PRIVATE_CHALLENGE_STATE` (required)
- ❑ `PROXY_SG_DOMAIN` (optional) If specified, its value is prepended to the username and separated with a backslash.

Authentication_form

The initial form, `authentication_form`, looks similar to the following:

```
<HTML>
<HEAD>
<TITLE>Enter Proxy Credentials for Realm $(cs-realm)</TITLE>
</HEAD>
<BODY>
<H1>Enter Proxy Credentials for Realm $(cs-realm)</H1>
<P>Reason for challenge: $(exception.last_error)
<P>$(x-auth-challenge-string)
<FORM METHOD="POST" ACTION=$(x-cs-auth-form-action-url)
$(x-cs-auth-form-domain-field)
<P>Username: <INPUT NAME="PROXY_SG_USERNAME" MAXLENGTH="64"
VALUE=$(cs-username) ></P>
<P>Password: <INPUT TYPE="PASSWORD" NAME="PROXY_SG_PASSWORD"
MAXLENGTH="64"></P>
<INPUT TYPE="HIDDEN" NAME="PROXY_SG_REQUEST_ID" VALUE=$(x-cs-auth-
request-id)>
<INPUT TYPE="HIDDEN" NAME="PROXY_SG_PRIVATE_CHALLENGE_STATE"
VALUE=$(x-auth-private-challenge-state)>
```

```
<P><INPUT TYPE=SUBMIT VALUE="Submit"> <INPUT TYPE=RESET></P>
</FORM>
<P>$(exception.contact)
</BODY>
</HTML>
```

If the realm is an IWA realm, the `$(x-cs-auth-form-domain-field)` substitution expands to:

```
<P>Domain: <INPUT NAME=PROXY_SG_DOMAIN MAXLENGTH=64 VALUE=$(x-cs-auth-
domain)>
```

If you specify `$(x-cs-auth-form-domain-field)`, you do not need to explicitly add the domain input field.

For comparison, the `new_pin_form` and `query_form` look similar to the following:

New_pin_form

```
<HTML>
<HEAD>
<TITLE>Create New PIN for Realm $(cs-realm)</TITLE>
<SCRIPT LANGUAGE="JavaScript"><!--
function validatePin() {
var info;
var pin = document.pin_form.PROXY_SG_PASSWORD;
if (pin.value != document.pin_form.PROXY_SG_RETYPE_PIN.value) {
    info = "The PINs did not match. Please enter them again.";
} else {
    // Edit this regular expression to match local PIN
    definition
    var re=/^[A-Za-z0-9]{4,16}$/
    var match=re.exec(pin.value);
    if (match == null) {
        info = "The PIN must be 4 to 16 alphanumeric
characters";
    } else {
        return true;
    }
}
alert(info);
pin.select();
pin.focus();
return false;
} // -->
</script>
</HEAD>
<BODY>
<H1>Create New PIN for Realm $(cs-realm)</H1>
<P>$(x-auth-challenge-string)
<FORM NAME="pin_form" METHOD="POST" ACTION=$(x-cs-auth-form-action-
url)ONSUBMIT="return validatePin()">
$(x-cs-auth-form-domain-field)
<P> Enter New Pin: <INPUT TYPE=PASSWORD NAME="PROXY_SG_PASSWORD"
MAXLENGTH="64"></P>
<P>Retype New Pin: <INPUT TYPE=PASSWORD NAME="PROXY_SG_RETYPE_PIN"
MAXLENGTH="64"></P>
<INPUT TYPE=HIDDEN NAME="PROXY_SG_USERNAME" VALUE=$(cs-username)>
```

```
<INPUT TYPE=HIDDEN NAME="PROXY_SG_REQUEST_ID" VALUE=$(x-cs-auth-
request-id)>
<INPUT TYPE=HIDDEN NAME="PROXY_SG_PRIVATE_CHALLENGE_STATE" VALUE=$(x-
auth-private-challenge-state)>
<P><INPUT TYPE=SUBMIT VALUE="Submit"></P>
</FORM>
<P>$(exception.contact)
</BODY>
</HTML>
```

Query_form

```
<HTML>
<HEAD>
<TITLE>Query for Realm $(cs-realm)</TITLE>
</HEAD>
<BODY>
<H1>Query for Realm $(cs-realm)</H1>
<P>$(x-auth-challenge-string)
<FORM METHOD="POST" ACTION=$(x-cs-auth-form-action-url)>
$(x-cs-auth-form-domain-field)
<INPUT TYPE=HIDDEN NAME="PROXY_SG_USERNAME" VALUE=$(cs-username)>
<INPUT TYPE=HIDDEN NAME="PROXY_SG_REQUEST_ID" VALUE=$(x-cs-auth-
request-id)>
<INPUT TYPE=HIDDEN NAME="PROXY_SG_PRIVATE_CHALLENGE_STATE" VALUE=$(x-
auth-private-challenge-state)>
<INPUT TYPE=HIDDEN NAME="PROXY_SG_PASSWORD">
<P><INPUT TYPE=SUBMIT VALUE="Yes"
ONCLICK="PROXY_SG_PASSWORD.value='Y'">
<INPUT TYPE=SUBMIT VALUE="No" ONCLICK="PROXY_SG_PASSWORD.value='N'"></
P>
</FORM>
<P>$(exception.contact)
</BODY>
</HTML>
```

User/Realm CPL Substitutions for Authentication Forms

CPL user/realm substitutions that are common in authentication form exceptions are listed below. The syntax for a CPL substitution is:

\$(CPL_substitution)

group	user-name	x-cs-auth-request-id
groups	user.x509.issuer	x-cs-auth-domain
realm	user.x509.serialNumber	x-cs-auth-form-domain-field
user	user.x509.subject	x-cs-auth-form-action-url
cs-realm	x-cs-auth-request-id	x-auth-challenge-string
	x-auth-private-challenge- state	

Note: Any substitutions that are valid in CPL and in other exceptions are valid in authentication form exceptions. There is no realm restriction on the number of authentication form exceptions you can create. You can have an unlimited number of forms, but make them as generic as possible to cut down on maintenance.

For a discussion of CPL and a complete list of CPL substitutions, as well as a description of each substitution, refer to the *Content Policy Language Guide*.

Storage Options

When a request requiring the user to be challenged with a form contains a body, the request is stored on the ProxySG appliance while the user is being authenticated. Storage options include:

- ❑ the maximum request size
- ❑ the expiration of the request
- ❑ whether to verify the IP address of the client requesting against the original request
- ❑ whether to allow redirects from the origin server

The storage options are global, applying to all form exceptions you use.

The global allow redirects configuration option can be overridden on a finer granularity in policy using the `authenticate.redirect_stored_requests(yes|no)` action.

Configuring Forms-Based Authentication

To create and put into use forms-based authentication, you must complete the following steps:

- ❑ Create a new form or edit one of the existing authentication form exceptions. See "[Creating a New Form](#)" on page 1205.
- ❑ Set storage options. See "[Storage Options](#)" on page 1204.
- ❑ Set policies. See "[Using CPL with Forms-Based Authentication](#)" on page 1207.

Creating and Editing a Form

You can create a new form or you can edit one of the existing ones as described in the following sections:


- ❑ "[Creating a New Form](#)" on page 1205
- ❑ "[Editing an Existing Form](#)" on page 1205

Creating a New Form

When you create a new form, you must define its type (**authentication_form**, **new_pin_form**, or **query_form**). The form is created from the default definition for that type.

To create an authentication form:

1. Select the **Configuration > Authentication > Forms > Authentication Forms** tab.
2. Select **New** to create a new form. The Add list item dialog displays.


A screenshot of the 'Add list item' dialog box. The dialog has a blue title bar with the text 'Add list item' and a red close button. Inside the dialog, there is a section titled 'Add Authentication Form'. Below this title, there are two input fields: 'Name:' followed by a text box, and 'Type:' followed by a dropdown menu showing 'Authentication Form'. At the bottom of the dialog, there are two buttons: 'OK' and 'Cancel'.

3. Enter the form **Name**.
 4. From the Type drop-down list, select a authentication form type. If you do not know the difference, see ["About Authentication Forms"](#) on page 1199.
- ☐ Click **OK**.

Editing an Existing Form

To edit a form:

1. Select **Configuration > Authentication > Forms**.
2. Select the form you want to edit and click **Edit**. The Edit Authentication Form dialog box is displayed.

A screenshot of the 'Edit Authentication Form new_pin_form' dialog box. The dialog has a blue title bar with the text 'Edit Authentication Form new_pin_form' and a red close button. Inside the dialog, there is a section titled 'Edit Authentication Form new_pin_form'. Below this title, there is a label 'Install Authentication Form from:' followed by a dropdown menu showing 'Remote URL' and an 'Install' button. At the bottom of the dialog, there are two buttons: 'OK' and 'Cancel'.

Note: **View** in the Authentication Forms panel and **View** in the Default Definitions panel have different functions. **View** in the Authentication Forms panel allows you to view the form you highlighted; **View** in the Default Definitions panel allows you view the original, default settings for each form. This is important in an upgrade scenario; any forms already installed will not be changed. You can compare existing forms to the default version and decide if your forms need to be modified.

3. Select one of the following installation options from the **Install Authentication Form from** drop-down list:
 - **Remote URL**—Enter the fully-qualified URL, including the filename, where the authentication form is located. To view the file before installing it, click **View**. Click **Install**. To view the results, click **Results**; to close the dialog when through, click **OK**.
 - **Local File**—Click **Browse** to bring up the Local File Browse window. Browse for the file on the local system. Open it and click **Install**. When the installation is complete, a results window opens. View the results; to close the window, click **Close**.
 - **Text Editor**—The current authentication form is displayed in the text editor. You can edit the form in place.
4. To install the form, click **Install**. When the installation is complete, a results window opens.

Setting Storage Options

This section discusses how to set storage options for authentication forms. For more information, see ["Storage Options"](#) on page 1204.

To set storage options:

1. Select the **Configuration > Authentication > Forms > Request Storage** tab.

The screenshot shows the 'Request Storage' configuration window. It has two tabs: 'Authentication Forms' and 'Request Storage'. The 'Request Storage' tab is active. Inside the window, there are four settings:

- 2 points to the 'Request Storage' tab header.
- 2 points to the 'Maximum request size to store (Megabytes):' field, which contains the value '50'.
- 3 points to the 'Request object expiry time (seconds):' field, which contains the value '300'.
- 4 points to the 'Verify the IP address against the original request' checkbox, which is checked.
- 5 points to the 'Allow redirects' checkbox, which is unchecked.

2. In the **Maximum request size to store (Megabytes)** field, enter the maximum POST request size allowed during authentication. The default is 50 megabytes.

3. In the **Request object expiry time (seconds)** field, enter the amount of time before the stored request expires. The default is 300 seconds (five minutes). The expiry time should be long enough for the user to fill out and submit the authentication form.
4. If you do not want the ProxySG to **Verify the IP address against the original request**, deselect that option. The default is to verify the IP address.
5. To **Allow redirects** from the origin servers, select the check box. The default is to not allow redirects from origin servers. Enable this option if you know that the redirects are going to a known server.

Note: During authentication, the user's POST is redirected to a GET request. The client therefore automatically follows redirects from the origin server. Because the ProxySG is converting the GET to a POST and adding the post data to the request before contacting the origin server, the administrator must explicitly specify that redirects to these POSTs requests can be automatically followed.

6. Click **Apply**.

Using CPL with Forms-Based Authentication

To use forms-based authentication, you must create policies that enable it and also control which form is used in which situations. A form must exist before it can be referenced in policy.

- ❑ Which form to use during authentication is specified in policy using one of the CPL conditions `authenticate.form(form_name)`, `authenticate.new_pin_form(form_name)`, or `authenticate.query_form(form_name)`.

These conditions override the use of the initial forms for the cases where a new pin form needs to be displayed or a query form needs to be displayed. All three of the conditions verify that the form name has the correct type.

Note: Each of these conditions can be used with the form authentication modes only. If no form is specified, the form defaults to the CPL condition for that form. That is, if no name is specified for `authenticate.form(form_name)`, the default is `authentication_form`; if no name is specified for `authenticate.new_pin_form(form_name)`, the default is `authenticate.new_pin_form`, and if no name is specified for `authenticate.query_form(form_name)`, the default is `authenticate.query_form`.

- ❑ Using the `authentication.mode( )` property selects a combination of challenge type and surrogate credentials. The `authentication.mode( )` property offers several options specifically for forms-based authentication:

- **Form-IP**—The user's IP address is used as a surrogate credential. The form is presented whenever the user's credential cache entry expires.
 - **Form-Cookie**—Cookies are used as surrogate credentials. The cookies are set on the OCS domain only, and the user is presented with the form for each new domain. This mode is most useful in reverse proxy scenarios where there are a limited number of domains.
 - **Form-Cookie-Redirect**—The user is redirected to the authentication virtual URL before the form is presented. The authentication cookie is set on both the virtual URL and the OCS domain. The user is only challenged when the credential cache entry expires.
 - **Form-IP-redirect** —This is similar to **Form-IP** except that the user is redirected to the authentication virtual URL before the form is presented.
- ❑ If you authenticate users who have third-party cookies explicitly disabled, you can use the `authenticate.use_url_cookie( )` property.
 - ❑ Since the `authentication.mode( )` property is defined as a form mode (above) in policy, you do not need to adjust the default authenticate mode through the CLI.
 - ❑ Using the `authenticate.redirect_stored_requests(yes|no)` action allows granularity in policy over the global allow redirect config option.

For information on using these CPL conditions and properties, refer to *Content Policy Language Guide*.

Troubleshooting Forms-Based Authentication

- ❑ If the user is supposed to be challenged with a form on a request for an image or video, the ProxySG returns a 403 error page instead of the form. If the reason for the challenge is that the user's credentials have expired and the object is from the same domain as the container page, then reloading the container page results in the user receiving the authentication form and being able to authenticate. However, if the client browser loads the container page using an existing authenticated connection, the user might still not receive the authentication form.

Closing and reopening the browser should fix the issue. Requesting a different site might also cause the browser to open a new connection and the user is returned the authentication form.

If the container page and embedded objects have a different domain though and the authentication mode is **form-cookie**, reloading or closing and reopening the browser might not fix the issue, as the user is never returned a cookie for the domain the object belongs to. In these scenarios, Blue Coat recommends that policy be written to either bypass authentication for that domain or to use a different authentication mode such as **form-cookie-redirect** for that domain.

- ❑ Forms-based authentication works with Web browsers only.

- ❑ Because forms only support Basic authentication, authentication-form exceptions cannot be used with a Certificate realm. If a form is in use and the authentication realm is or a Certificate realm, you receive a configuration error.
- ❑ User credentials are sent in plain text. However, they can be sent securely using SSL if the virtual URL is HTTPS.
- ❑ Because not all user requests support forms (such as WebDAV and streaming), create policy to bypass authentication or use a different authentication mode with the same realm for those requests.

Chapter 66: Authentication and Authorization Errors

Following is the list of all groups and individual errors that can be permitted during authentication and authorization. The first table lists the groups and the individual errors within each group. The second table lists all of the individual errors along with descriptions of the errors.

Table 66–1 Groups and Individual Errors

Error Group	CPL	Members	Description
All	All	account_disabled account_expired account_locked_out account_must_change_password account_restricted account_wrong_place account_wrong_time agent_config_changed agent_config_cmd_failed agent_connection_failed agent_init_failed agent_no_groups_provided agent_resource_not_protected agent_too_many_retries agent_unsupported_scheme authorization_username_too_long basic_password_too_long basic_username_too_long cannot_decrypt_secret cannot_determine_authorization_username cannot_determine_full_username cannot_determine_username cannot_expand_credentials_substitution cannot_redirect_connect cannot_redirect_https_to_http cannot_setup_working_dir cert_explicit_unsupported certificate_missing credential_decode_failure credentials_mismatch	Includes all errors that can be permitted. If this group includes errors, such as need_credentials. If permitted, these errors result in the user never being challenged. As this is not the desired behavior for most realms (for example, the user should be given the chance to enter credentials) do not permit this group when using challenge realms. Instead, use combinations of the other error groups as appropriate.

Table 66–1 Groups and Individual Errors (Continued)

Error Group	CPL	Members	Description
		domain_controller_query_disabled expired_credentials form_does_not_support_connect form_requires_basic_support general_authentication_error general_authorization_error guest_user invalid_ip invalid_license invalid_local_user_list invalid_realm invalid_search_credentials invalid_surrogate issuer_too_long ldap_busy ldap_filter_error ldap_inappropriate_auth ldap_insufficient_access ldap_invalid_credentials ldap_invalid_dn_syntax ldap_loop_detect ldap_no_such_attribute ldap_no_such_object ldap_partial_results ldap_server_down ldap_timelimit_exceeded ldap_timeout ldap_unavailable ldap_unwilling_to_perform missing_base_dn missing_form_configuration multiple_users_matched need_credentials netbios_failure netbios_cannot_send netbios_multiple_users netbios_no_computer_name netbios_no_domain_name netbios_no_user_name netbios_recv_failed netbios_reply_invalid netbios_reply_timeout no_offbox_url_specified no_servers no_user_in_cert not_attempted not_ssl offbox_abort offbox_missing_secret offbox_process_create_failed offbox_protocol_error offbox_server_down	

Table 66–1 Groups and Individual Errors (Continued)

Error Group	CPL	Members	Description
		offbox_server_unreachable offbox_timeout otp_already_used password_too_long radius_socket_interface rdns_cannot_determine_name rdns_failed redirect_from_vh sspi_context_lost sspi_context_too_old sspi_domain_controller_not_found sspi_invalid_handle sspi_invalid_mechanism sspi_invalid_token sspi_invalid_type3_message sspi_logon_denied sspi_logon_type_not_granted sspi_no_authenticating_authority sspi_null_lm_password sspi_process_create_failed sspi_rpc_error sspi_service_disabled sspi_timeout sspi_unable_to_connect_to_agent subject_too_long too_many_users unable_to_query_client unknown_user user_domain_not_trusted username_too_long	
Communication Error	communication_error	agent_connection_failed ldap_busy ldap_loop_detect ldap_server_down ldap_unavailable ldap_unwilling_to_perform netbios_cannot_send netbios_reply_invalid no_servers radius_socket_interface sspi_no_authenticating_authority sspi_rpc_error sspi_unable_to_connect_to_agent	Includes communication errors with BCAAA, LDAP, and RADIUS servers and during NetBIOS queries.
Configuration Changed	configuration_change	agent_config_changed offbox_abort	The ProxySG has been notified that configuration affecting the realm has been changed off-box. Used primarily with SiteMinder and COREid realms.

Table 66–1 Groups and Individual Errors (Continued)

Error Group	CPL	Members	Description
General Authentication Failure	general_authentication_failure	general_authentication_error	A general authentication error has occurred. This is returned when a specific error does not apply. It does not include all authentication errors.
General Authorization Failure	general_authorization_failure	cannot_determine_authorization_username general_authorization_error	A general authorization error has occurred. This is returned when a specific error does not apply. It does not include all authorization errors. This can be returned as an authentication error in realms that do not support specifying a separate authorization realm.

Table 66–1 Groups and Individual Errors (Continued)

Error Group	CPL	Members	Description
General Offbox Error	offbox_error	agent_connection_failed agent_init_failed cannot_determine_full_username ldap_busy ldap_loop_detect ldap_server_down ldap_timelimit_exceeded ldap_timeout ldap_unavailable ldap_unwilling_to_perform netbios_failure netbios_cannot_send netbios_multiple_users netbios_no_computer_name netbios_no_domain_name netbios_no_user_name netbios_recv_failed netbios_reply_invalid netbios_reply_timeout no_servers offbox_process_create_failed offbox_protocol_error offbox_server_down offbox_server_unreachable offbox_timeout radius_socket_interface rdns_cannot_determine_name rdns_failed sspi_context_lost sspi_context_too_old sspi_invalid_mechanism sspi_no_authenticating_authority sspi_process_create_failed sspi_rpc_error sspi_timeout sspi_unable_to_connect_to_agent	Includes all errors that can result with failures found during any offbox configuration or communications. It includes all errors found in the Communication Error group.

Table 66–1 Groups and Individual Errors (Continued)

Error Group	CPL	Members	Description
General Onbox Error	onbox_error	onbox_BASE64_decode_failure onbox_BASE64_encode_failure onbox_clock_skew onbox_create_domain_trust_failed onbox_create_refresher_thread_failed onbox_domain_join_error onbox_domain_not_found onbox_domain_offline onbox_gss_error onbox_gss_unable_to_export_username onbox_gss_unable_to_retrieve_pac onbox_krb5_error onbox_sid_info_not_available onbox_unmapped_error onbox_username_wrong_format onbox_user_not_found onbox_wrong_service_principal	Errors found during configuration or communication with an onbox authentication realm, such as IWA Direct.
Ident Error	ident_error		Errors found during Ident query
Initialization Error	initialization_error	agent_init_failed offbox_process_create_failed sspi_process_create_failed	Errors related to initializing the realm.
Internal Error	internal_error		Any internal error.
Invalid BCAAA Request	invalid_bcaaa_request	sspi_context_lost sspi_context_too_old sspi_invalid_mechanism	Includes errors returned if the request sent to BCAAA is invalid. Applies to IWA realms only.

Table 66–1 Groups and Individual Errors (Continued)

Error Group	CPL	Members	Description
Invalid Configuration	invalid_configuration	agent_config_cmd_failed agent_no_groups_provided agent_resource_not_protected agent_too_many_retries agent_unsupported_scheme cannot_decrypt_secret cannot_determine_full_username cannot_determine_username cannot_setup_working_dir cert_explicit_unsupported domain_controller_query_disabled form_does_not_support_connect form_requires_basic_support invalid_local_user_list invalid_realm invalid_search_credentials ldap_filter_error ldap_inappropriate_auth ldap_insufficient_access ldap_invalid_dn_syntax ldap_no_such_attribute ldap_no_such_object ldap_partial_results missing_base_dn missing_form_configuration no_offbox_url_specified no_servers not_ssl offbox_missing_secret offbox_protocol_error offbox_server_unreachable sspi_domain_controller_not_found sspi_logon_type_not_granted sspi_null_lm_password sspi_service_disabled	Includes any errors that resulted from a possible misconfiguration of the ProxySG. These errors usually require administrator action to address.
Invalid License	invalid_license	invalid_license	An invalid license was found for an authentication component.
Invalid NetBIOS Reply	invalid_netbios_reply	netbios_failure netbios_multiple_users netbios_no_computer_name netbios_no_domain_name netbios_no_user_name netbios_recv_failed	The NetBIOS reply was invalid.

Table 66–1 Groups and Individual Errors (Continued)

Error Group	CPL	Members	Description
Invalid User Information	invalid_user_info	authorization_username_too_long basic_password_too_long basic_username_too_long cannot_expand_credentials_substitution credential_decode_failure credentials_mismatch general_authentication_error invalid_surrogate issuer_too_long ldap_invalid_credentials otp_already_used password_too_long sspi_invalid_handle sspi_invalid_token sspi_invalid_type3_message sspi_logon_denied subject_too_long user_domain_not_trusted username_too_long	Includes errors that result from invalid user information being entered.
RDNS Failure	rdns_failure	rdns_cannot_determine_name rdns_failed	Errors found during Reverse DNS lookup.
Redirect Error	redirect_error	cannot_redirect_connect cannot_redirect_https_to_http redirect_from_vh	Errors found while attempting to redirect the user's request for authentication. Only returned when using a redirect authentication mode.
Request Timeout	request_timeout	ldap_timelimit_exceeded ldap_timeout netbios_reply_timeout offbox_timeout sspi_timeout	Includes timeout errors with authentication servers.
Single Sign-on Failure	sso_failure	invalid_ip multiple_users_matched too_many_users unknown_user unable_to_query_client	Errors returned during Single Sign-on authentication. These errors apply to Windows SSO and Novell SSO realms only.

Table 66–1 Groups and Individual Errors (Continued)

Error Group	CPL	Members	Description
User Account Error	user_account_error	account_disabled account_expired account_locked_out account_must_change_password account_restricted account_wrong_place account_wrong_time expired_credentials	Errors with the user's account.
User Credentials Required	credentials_required	certificate_missing guest_user need_credentials no_user_in_cert	User credentials are required. Do not permit this error if the user should be challenged for credentials.

Table 66–2 Individual Errors

Error Name	Description	Groups
account_disabled	Account is disabled.	All User Account Error
account_expired	Account has expired.	All User Account Error
account_locked_out	Account is locked out.	All User Account Error
account_must_change_password	Account password must be changed.	All User Account Error
account_restricted	Account is restricted.	All User Account Error
account_wrong_place	Account cannot be used from this location.	All User Account Error
account_wrong_time	Account logon time restricted - cannot be used now.	All User Account Error
agent_config_changed	Agent reports server configuration has changed; please try your request again.	All Configuration Changed
agent_config_cmd_failed	Configuration of the authentication agent failed	All Invalid Configuration
agent_connection_failed	The authentication agent could not communicate with its authority.	All Communication Error General Off-box Error
agent_init_failed	The authentication agent failed to initialize.	All Initialization Error General Off-box Error

Table 66–2 Individual Errors (Continued)

Error Name	Description	Groups
agent_no_groups_provided	The authentication agent did not receive the group list from the server.	All Invalid Configuration
agent_resource_not_protected	The authentication agent reports that the resource is not protected.	All Invalid Configuration
agent_too_many_retries	Agent configuration failed.	All Invalid Configuration
agent_unsupported_scheme	The requested authentication scheme is not supported.	All Invalid Configuration
authorization_username_too_long	The resolved authorization username is too long.	All Invalid User Information
basic_password_too_long	Basic password is too long.	All Invalid User Information
basic_username_too_long	Basic username is too long.	All Invalid User Information
cannot_decrypt_secret	Cannot decrypt shared secret.	All Invalid Configuration
cannot_determine_authorization_username	Could not determine the authorization username.	All General Authorization Failure
cannot_determine_full_username	Could not determine full user name.	All Invalid Configuration General Off-box Error
cannot_determine_username	Agent could not determine simple user name.	All Invalid Configuration
cannot_expand_credentials_substitution	The substitution used to determine the credentials could not be expanded.	All Invalid User Information
cannot_redirect_connect	Cannot use origin-redirect or form-redirect for CONNECT method (explicit proxy of https URL)	All Redirect Error
cannot_redirect_https_to_http	Cannot redirect an HTTPS request to an HTTP virtual URL	All Redirect Error
cannot_setup_working_dir	Unable to setup working directory for COREid AccessGate	All Invalid Configuration
cert_explicit_unsupported	Certificate authentication not supported for explicit proxy.	All Invalid Configuration

Table 66–2 Individual Errors (Continued)

Error Name	Description	Groups
certificate_missing	No certificate found. Check that verify-client is set on https service.	All User Credentials Required
credential_decode_failure	Unable to decode base64 credentials.	All Invalid User Information
credentials_mismatch	Credentials did not match.	All Invalid User Information
domain_controller_query_disabled	Windows SSO Domain Controller Querying is not enabled on the Single Sign-on agent.	All Invalid Configuration
expired_credentials	Credentials on back-end server have expired.	All User Account Error
form_does_not_support_connect	Cannot use form authentication for CONNECT method (explicit proxy of https URL)	All Invalid Configuration
form_requires_basic_support	Form authentication requires the realm to support Basic credentials.	All Invalid Configuration
general_authentication_error	General authentication failure due to bad user ID or authentication token.	All General Authentication Failure Invalid User Information
general_authorization_error	Unable to authorize authenticated user.	All General Authorization Failure
guest_user	Credentials required.	All User Credentials Required
invalid_ip	The IP address of this computer could not be determined by the Single Sign-on agent.	All Single Sign-on Failure
invalid_license	The license for the configured realm does not exist or is invalid. A valid license must be installed.	All Invalid License
invalid_local_user_list	Invalid local user list.	All Invalid Configuration
invalid_realm	The specified realm is invalid.	All Invalid Configuration

Table 66–2 Individual Errors (Continued)

Error Name	Description	Groups
invalid_search_credentials	The LDAP search credentials are invalid.	All Invalid Configuration
invalid_surrogate	The surrogate is invalid for the specified realm.	All Invalid User Information
issuer_too_long	Certificate's issuer string is too long.	All Invalid User Information
ldap_busy	LDAP: server busy.	All Communication Error General Off-box Error
ldap_filter_error	LDAP: filter error.	All Invalid Configuration
ldap_inappropriate_auth	LDAP: inappropriate authentication.	All Invalid Configuration
ldap_insufficient_access	LDAP: insufficient access.	All Invalid Configuration
ldap_invalid_credentials	LDAP: invalid credentials.	All Invalid User Information
ldap_invalid_dn_syntax	LDAP: invalid DN syntax.	All Invalid Configuration
ldap_loop_detect	LDAP: loop detected.	All Communication Error General Off-box Error
ldap_no_such_attribute	LDAP: No such attribute.	All Invalid Configuration
ldap_no_such_object	LDAP: no such object.	All Invalid Configuration
ldap_partial_results	LDAP server returned partial results.	All Invalid Configuration
ldap_server_down	Could not connect to LDAP server.	All Communication Error General Off-box Error
ldap_timelimit_exceeded	LDAP server exceeded time limit.	All Request Timeout General Off-box Error

Table 66–2 Individual Errors (Continued)

Error Name	Description	Groups
ldap_timeout	The LDAP request timed out.	All Request Timeout General Off-box Error
ldap_unavailable	LDAP: service unavailable.	All Communication Error General Off-box Error
ldap_unwilling_to_perform	LDAP: server unwilling to perform requested action.	All Communication Error General Off-box Error
missing_base_dn	No base DN's are configured.	All Invalid Configuration
missing_form_configuration	Form authentication is not properly configured	All Invalid Configuration
multiple_users_matched	The user query resulted in multiple users. A unique user could not be determined.	All Single Sign-on Failure
need_credentials	Credentials are missing.	All User Credentials Required
netbios_failure	NetBIOS reply did not contain data needed for authentication.	All Invalid NetBIOS Reply General Off-box Error
netbios_cannot_send	Could not send NetBIOS query.	All Communication Error General Off-box Error
netbios_multiple_users	NetBIOS reply contained multiple user names.	All Invalid NetBIOS Reply General Off-box Error
netbios_no_computer_name	Could not determine computer name from NetBIOS reply.	All Invalid NetBIOS Reply General Offbox Error
netbios_no_domain_name	Could not determine domain name from NetBIOS reply.	All Invalid NetBIOS Reply General Off-box Error

Table 66–2 Individual Errors (Continued)

Error Name	Description	Groups
netbios_no_user_name	NetBIOS reply did not contain the username.	All Invalid NetBIOS Reply General Off-box Error
netbios_recv_failed	Failed to receive reply to NetBIOS query.	All Invalid NetBIOS Reply General Off-box Error
netbios_reply_invalid	Reply to NetBIOS query was invalid.	All Communication Error General Off-box Error
netbios_reply_timeout	Timed out awaiting reply to NetBIOS query.	All Request Timeout General Off-box Error
no_offbox_url_specified	Off-box redirects are configured but no off-box URL is specified.	All Invalid Configuration
no_servers	No usable authentication servers found.	All Communication Error Invalid Configuration General Off-box Error
no_user_in_cert	Could not retrieve username from certificate.	All User Credentials Required
none	Status successful.	
not_attempted	The method has not been attempted.	All
not_ssl	SSL is required but connection is not using it (check virtual-url).	All Invalid Configuration
offbox_abort	The request was aborted due to a change in configuration.	All Configuration Changed
offbox_missing_secret	Secret is not defined for authentication realm	All Invalid Configuration
offbox_process_create_failed	Could not create offbox authentication processes	All Initialization Error General Off-box Error
offbox_protocol_error	The authentication server returned an invalid result.	All Invalid Configuration General Off-box Error

Table 66–2 Individual Errors (Continued)

Error Name	Description	Groups
offbox_server_down	The authentication server cannot process requests.	All General Off-box Error
offbox_server_unreachable	The authentication server could not be contacted.	All Invalid Configuration General Off-box Error
offbox_timeout	The request timed out while trying to authenticate. The authentication server may be busy or offline.	All Request Timeout General Off-box Error
otp_already_used	The one-time password has already been used	All Invalid User Information
password_too_long	Password is too long.	All Invalid User Information
radius_socket_interface	RADIUS received an unexpected socket error.	All Communication Error General Off-box Error
rdns_cannot_determine_name	Could not determine user name from client host name.	All RDNS Failure General Off-box Error
rdns_failed	Reverse DNS address resolution failed.	All RDNS Failure General Off-box Error
redirect_from_vh	Redirecting from the virtual host.	All Redirect Error
sspi_context_lost	Authentication agent rejected request (context lost).	All Invalid BCAA Request General Off-box Error
sspi_context_too_old	Authentication agent rejected request - too old.	All Invalid BCAA Request General Off-box Error
sspi_domain_controller_not_found	Cannot find domain controller.	All Invalid Configuration
sspi_invalid_handle	SSPI protocol error - invalid context handle.	All Invalid User Information

Table 66–2 Individual Errors (Continued)

Error Name	Description	Groups
sspi_invalid_mechanism	Authentication agent rejected request - Invalid mechanism requested.	All Invalid BCAA Request General Off-box Error
sspi_invalid_token	The credentials provided are invalid.	All Invalid User Information
sspi_invalid_type3_message	Client sent invalid NTLM Type 3 message.	All Invalid User Information
sspi_logon_denied	The logon failed.	All Invalid User Information
sspi_logon_type_not_granted	Requested logon type not granted.	All Invalid Configuration
sspi_no_authenticating_authority	No authority could be contacted for authentication.	All Communication Error General Off-box Error
sspi_null_lm_password	Windows NT password too complex for LanMan.	All Invalid Configuration
sspi_process_create_failed	NTLM realm could not create administrator processes.	All Initialization Error General Off-box Error
sspi_rpc_error	Connection to authentication agent lost.	All Communication Error General Off-box Error
sspi_service_disabled	SSPI service disabled.	All Invalid Configuration
sspi_timeout	Authentication agent did not respond to request in time.	All Request Timeout General Off-box Error
sspi_unable_to_connect_to_agent	Unable to connect to authentication agent.	All Communication Error General Off-box Error
subject_too_long	Certificate's subject string is too long.	All Invalid User Information

Table 66–2 Individual Errors (Continued)

Error Name	Description	Groups
too_many_users	More than one user is logged onto this computer. Only one user can be logged on for Single Sign-on authentication.	All Single Sign-on Failure
unable_to_query_client	The client workstation could not be queried by the Single Sign-on agent.	All Single Sign-on Failure
unknown_user	The user could not be determined by the Single Sign-on agent.	All Single Sign-on Failure
user_domain_not_trusted	The specified domain is not trusted.	All Invalid User Information
username_too_long	Specified username is too long.	All Invalid User Information

Chapter 67: Configuring Adapters and Virtual LANs

This section describes ProxySG appliance network adapters, the adapter interfaces, and how to configure the ProxySG appliance to function within a Virtual LAN (VLAN) environment. Although you most likely have performed initial configuration tasks to get the ProxySG appliance live on the network, this section provides additional conceptual information to ensure the configuration matches the deployment requirement.

Topics in this Section

The following topics are covered in this section:

- ❑ ["How Do I...?"](#) on page 1229—Begin here if you are not sure of the answer you seek.
- ❑ ["How ProxySG appliance Adapters Interact on the Network"](#) on page 1230
- ❑ ["About VLAN Configurations"](#) on page 1233
- ❑ ["Changing the Default Adapter and Interface Settings"](#) on page 1237
- ❑ ["Viewing Interface Statistics"](#) on page 1248
- ❑ ["Detecting Network Adapter Faults"](#) on page 1249

How Do I...?

Identify the task to perform and click the link:

How do I...?	See...
Verify the ProxySG appliance is connected properly based on the basic deployment type, such as bridging and in-path?	"About WAN and LAN Interfaces" on page 1230
Learn basic information about virtual LAN (VLAN) deployments?	"About VLAN Configurations" on page 1233
Change the settings for default link speeds for interfaces?	"About Link Settings" on page 1232

How do I...?	See...
Verify that traffic is flowing through the interfaces and see what type of traffic it is?	"Viewing Interface Statistics" on page 1248
Troubleshoot interface connectivity?	"Detecting Network Adapter Faults" on page 1249

How ProxySG appliance Adapters Interact on the Network

Each ProxySG appliance ships with one or more network adapters installed on the system, each with one or more interfaces (the number of available interfaces varies by ProxySG appliance model).

Note: In Blue Coat documentation, the convention for the interface is *adapter:interface*. For example, 0:0.

About WAN and LAN Interfaces

Recent ProxySG appliance models have labels next to the physical interfaces (on the appliance backplate) that identify the WAN and LAN links. These interface labels are also hard coded in SGOS 5.3.x and later and are displayed in the respective interface graphics in the Management Console. Based on your deployment type (the ProxySG appliance directly in-path between users and a router or the ProxySG appliance connected to a router that resides in-path, virtually in-path, and explicit), verify the following connections:

- ❑ The ProxySG appliance is deployed in-path with bridging.




Figure 67–1 Connecting WAN and LAN interfaces in-path with bridging.

- ❑ Clients and WAN links connect to the ProxySG appliance transparently through a router with WCCP.




Figure 67–2 Connecting the LAN interface to a router with WCCP.

About Interception Options

The ProxySG appliance allows you to execute one of three actions upon intercepting traffic on a per-interface basis:

- ❑ **Allow:** Bridge/forward traffic and intercept appropriate traffic as defined by Proxy Services.
- ❑ **Bypass:** Bridge/forward all traffic without interception.
- ❑ **Firewall:** Drop (silently block) any traffic not related to established ProxySG appliance connections.

The following table describes what effect each allow-intercept option setting has on different traffic types.

Table 67–1 How each interception option affects connections.

Option	ProxySG appliance Settings		ProxySG appliance Management and Console Connections	Explicit Proxy Service Traffic	Transparent Proxy Service Traffic	Other Traffic
	reject-inbound	allow-intercept				
<i>Allow</i>	Disabled	Enabled	Intercepted	Intercepted	Intercepted	Forwarded
<i>Bypass</i>	Disabled	Disabled	Intercepted	Intercepted	Forwarded	Forwarded
<i>Firewall</i>	Enabled	Enabled/Disabled	Silently dropped	Silently dropped	Silently dropped	Silently dropped

The default intercept option depends on the type of license on this ProxySG appliance :

- ❑ Proxy Edition: The default is **Bypass transparent interception**.
- ❑ Mach 5 Edition: The default is **Allow transparent interception**. The ProxySG appliance performs normal proxy interception, as configured in **Configuration > Services**, for traffic on the interface. If you require this ProxySG appliance to perform interception of traffic on specific interface(s), set the other interfaces to either bypass (bridge/forward, but do not intercept traffic on it) or firewall it (drop all traffic not related to established proxy connections).

About Link Settings

By default, the ProxySG appliance auto-negotiates the interface *speed* and *duplex* settings with the switch or router to which it is connected.

- ❑ The ProxySG appliance supports multiple Ethernet modes. The speed setting is the maximum transfer speed, in Megabits or Gigabits per second (Mbps/Gbps), the interface supports.
- ❑ The duplex setting designates two-way traffic capabilities. In **Full** duplex mode, both devices may transmit to and from each other simultaneously, allowing each direction to use the maximum transfer speed without affecting the other direction. In **Half** duplex mode, only one device may transmit at any one time, effectively sharing the maximum transfer speed of the interface.

The ProxySG appliance's health monitoring capability provides alerts if interface use reaches warning and critical capacity levels. In **Full** duplex mode, the ProxySG appliance reports the larger percentage value of the sending and receiving values. For example, if the ProxySG appliance is receiving 20 Mbps and sending 40 Mbps on a 100 Mbps-capable interface, the reported value is 40%. If the same interface was set to half duplex, the reported value is 60%, or the aggregated values.

Blue Coat strongly recommends using the (default) auto-negotiation feature. The key issue is the ProxySG appliance settings must match the settings on the switch; therefore, if you manually change the settings on the ProxySG appliance, you must also match the settings on the router or switch.

Note: When the 100 Mbps Ethernet interfaces on the ProxySG appliance are connected to Gigabit Ethernet capable devices, they might incorrectly auto-negotiate when fail-open pass-through is used.

If both the interfaces on these ProxySG appliances are connected to Gigabit capable switches or hubs, Symantec recommends that you configure the link settings manually to 100 Mbps. To configure the link settings, see Step 3 in "[To configure a network adapter](#):" on page 1238.

The following table lists the results of various ProxySG appliance and router link settings for 100 Mbps speeds. The values are listed in the format: **speed/duplex**.

Table 67–2 Results for 100 Mbps link speed settings on the ProxySG appliance and the switch

Router/Switch Auto-negotiation Result (speed/duplex)	Router/Switch Interface Settings	ProxySG appliance Interface Setting	ProxySG appliance Auto-negotiation Result
100/Full Duplex	Auto	Auto	100/Full Duplex
N/A	100/Full Duplex	Auto	100/Half Duplex
N/A	100/Full Duplex	100/Full Duplex	N/A
100/Half Duplex	Auto	100/Full Duplex	N/A

The following table lists the results of various ProxySG appliance and router link settings for 1 Gbps speeds. The values are listed in the format: **speed/duplex**.

Table 67–3 Results for 1Gbps link speed settings on the ProxySG appliance and switch

Router/Switch Auto-negotiation Result	Router/Switch Interface Setting	ProxySG appliance Interface Setting	ProxySG appliance Auto-Negotiation Result
No Link	Auto	Gig/Full Duplex	No link
Gig/Full Duplex	Auto	Auto	Gig/Full Duplex
No Link	Gig/Full Duplex	Auto	No link

See Also

"Verifying the Health of Services Configured on the ProxySG" on page 1353.

About VLAN Configurations

Virtual LANs (VLANs) are *logical* network segments that allow hosts to communicate, regardless of physical network location. The benefit to this is that clients can be separated logically—based on organizational unit, for example—rather than based on physical connectivity to interfaces. The ProxySG appliance treats VLAN interfaces identically to traditional physical LAN interfaces.

VLAN segments are defined on the switch. The network administrator specifies which ports belong to which VLANs. The following diagram illustrates a port-based VLAN configuration. Clients on network segments attached to switch ports 1 and 2 belong to VLAN 1, which has the network address 10.0.1.x; network segments attached to switch ports 14 and 15 belong to VLAN 2, which has the network address 10.0.2.x.




Figure 67–3 Multiple VLANs connected to ports on one switch

As also illustrated in the diagram, clients of different OS types can reside within a VLAN. However, not all clients are able to detect (send or receive) VLAN-tagged packets.

About VLAN Trunking

Trunk ports are ports that carry traffic for more than one VLAN. They tag each packet with the VLAN ID in the packet header. Trunk ports are commonly used between switches and routers that must switch or route traffic from or to multiple VLANs. By default, VLAN trunking is enabled on the ProxySG appliance .

In the following diagram, multiple VLANs are connected by a trunk link between two switches.




Figure 67–4 Two switches connected by a trunk

About Native VLANs

Each switch port has a designated *native* VLAN. Traffic on the port associated with the native VLAN is not tagged. Traffic destined for VLANs other than the native VLAN is tagged.

The trunk link carries both the native VLAN and all other VLAN (tagged) packets, as illustrated in the following diagram.




Figure 67–5 A switch broadcasting native and regular VLAN traffic over a trunk

In this example, the client attached to port 7 belongs to VLAN 2. Even though port 7 is part of VLAN 2, it does not set tags or receive VLAN-tagged packets. The switch associates the traffic with VLAN 2 and tags it accordingly when appropriate. Conversely, it strips the VLAN 2 tag on the response. The trunk link carries VLAN 1 (the native) and 2 traffic to a router that forwards traffic for those VLANs.

Deployment complications arise when a device (other than a router) is required between switches. Any network device without VLAN-tagging support might drop or misinterpret the traffic.

As a best practice, do not deploy a device that is *not* configured to recognize VLAN-tagged traffic in-path of a trunk link.

Note: In Blue Coat documentation, the convention for VLAN is `adapter:interface.VLAN_ID`. Example: 1:0.10 refers the VLAN ID 10 on adapter 1, interface 0.

ProxySG appliance VLAN Support

The ProxySG appliance supports VLAN tagging and it is enabled by default; therefore, a ProxySG appliance can be deployed in-path with switches that are exchanging VLAN-tagged traffic. This allows for uninterrupted VLAN service, plus enables benefits gained with the proxy features.

The Management Console enables you to configure VLAN interfaces the same way you configure physical interfaces. After a VLAN is added, it appears in the list of network interfaces. Settings such as `allow-intercept` and `reject-inbound` are applicable to VLAN interfaces.

The most common deployment is a ProxySG appliance residing between two switches or a switch and a router; in these cases, preserving tagged packets is essential to proper network operation.




Figure 67-6 ProxySG appliance deployed between two switches

Based on this deployment:

- ❑ If configuration and policy allow, the ProxySG appliance accepts all packets regardless of their VLAN tag and passes them from one interface to the other with the original VLAN tag preserved.
- ❑ If a packet arrives on one interface tagged for VLAN 2, it remains on VLAN 2 when it is forwarded out on another interface. If a packet arrives untagged and the destination interface has a different native VLAN configured, the ProxySG appliance adds a tag to ensure the VLAN ID is preserved. Similarly, if a tagged packet arrives and the VLAN ID matches the native VLAN of the destination interface, the ProxySG appliance removes the tag before transmitting the packet.
- ❑ The ProxySG appliance strips the native VLAN tag on all outgoing traffic.

About Bridging and VLANs

On the ProxySG appliance, bridges can be created between two physical interfaces only. If you have configured virtual interfaces (VLANs), all the VLANs on the selected physical interfaces will be bridged.

Although VLANs are supported on bridges, the ProxySG appliance does not support creating a bridge group between VLANs when bridging or bypassing traffic. For example, when bridging you cannot send packets from VLAN 0:0.2 to 0:1.3.

Changing the Default Adapter and Interface Settings

The following procedure describes how to disable, enable, or change the default adapter and interface settings because of site-specific network requirements. These include inbound connection restrictions, link settings, browser/PAC file settings, and VLAN settings. Repeat the process if the system has additional adapters. By default:

- ❑ The ProxySG appliance allows the transparent interception of inbound connections.
- ❑ By default, the ProxySG appliance auto-negotiates link settings with the connected switch or router. Blue Coat recommends using auto-negotiation except under special circumstances.

Note: Rejecting inbound connections improperly or manually configuring link settings improperly might cause the ProxySG appliance to malfunction. Ensure that you know the correct settings before attempting either of these. If the ProxySG appliance fails to operate properly after changing these settings, contact Blue Coat Technical Support.

For more information, see one of the following topics:

- ❑ ["About Multiple IP Addresses"](#)
- ❑ ["Configuring a Network Adapter"](#) on page 1237

About Multiple IP Addresses

The ProxySG appliance allows you to bind multiple IP addresses to an interface, and typically, the assigned IP addresses are on the same subnet. Multiple IP addresses on an interface allows for managing one service under a specific IP and another service under a different IP. For example, you can assign one IP address for management services/console access and another IP address for managing proxy traffic. In addition, you could assign unique IP addresses to manage different services, that is have HTTP traffic on one and native FTP on another.

When using the ProxySG appliance in a mixed IPv4/IPv6 environment, you should assign IPv4 and IPv6 addresses to each interface. The IPv6 address can be link-local or global.

Configuring a Network Adapter



This section discusses how to configure a network adapter. For more information, see one of the following topics:

- ❑ ["Changing the Default Adapter and Interface Settings"](#) on page 1237
- ❑ ["About Multiple IP Addresses"](#) on page 1237

To configure a network adapter:

1. Select **Configuration > Network > Adapters > Adapters** tab.

Note: Different ProxySG appliance models have different adapter configurations.





2. Select the adapter and interface to configure:
 - a. In the **Physical Interfaces** area, select an adapter.
 - b. The **Link State** column displays the information in the form of:


Auto/Manual: *Speed* FDX/HDX

 - **Auto/Manual:** Whether or not the ProxySG appliance auto-negotiates with the router.
 - **Speed:** The maximum transfer speed available through the interface, depending on the type of Ethernet technology. The values are: **10**, **100**, and **1000** Mbps.


Note: An N/A status might indicate a network connectivity issue.

- **FDX/HDX:**
 - **FDX:** Full Duplex—the interface can simultaneously send and receive at the defined maximum speed (previous bullet). For example, a 100 Mbps full duplex link can send up to 100 Megabits per second (Mbps) of data and simultaneously receive up to 100 Mbps of data.
 - **HDX:** Half Duplex—the interface can only send data in one direction at a time. For example, a 100 Mbps half duplex link can only send and receive a *combined* maximum of 100 Mbps of data.
 - c. The **Aggregate Group** column displays aggregate group (or link) affiliation if applicable.
 - d. The **Bridge Group** column displays group affiliation. For more information about network bridging, see [Chapter 68: "Software and Hardware Bridges"](#) on page 1251.
3. To change a setting or name the interface, click **Configure interface #:#**. The Configure Interfaces dialog displays.


Dialog Area	Option
	Use Identification to associate ProxySG appliance interfaces with the connection purpose. For example, label an interface wan-sfodc to indicate a WAN-OP connection to a datacenter Concentrator in San Francisco.
 The default is Allow transparent interception. The ProxySG appliance performs normal proxy interception, as configured in Configuration > Services, for the traffic arriving on the interface. If you require this ProxySG appliance to perform interception on traffic from a specific interface or set of interfaces, set the other interfaces to either bypass the traffic (pass it through but not intercept it) or firewall it (block it completely). For more detailed information, see "About Interception Options" on page 1231.	Inbound connection options: • Allow transparent interception (default): The ProxySG appliance intercepts the appropriate traffic based on settings configured in Configuration > Services; all other traffic is bridged or forwarded. • Bypass transparent interception: The ProxySG appliance bridges or forwards <i>all</i> inbound traffic on this interface, regardless of the services configuration. • Firewall incoming traffic: The ProxySG appliance drops all inbound connections on this interface, regardless of the services configuration.

Dialog Area	Option
	Link settings: • Automatically sense link settings (default, recommended): The ProxySG appliance auto-negotiates the link settings for this interface. • Manually configure link settings: Select the options that meet your network requirements. This method requires a consistent configuration on the router or switch connected to this ProxySG appliance. Half is not available for an aggregate interface

4. Click **OK** to close the dialog.
5. Click **Apply** to save changes to the adapter/interface settings.
6. Next step:
 - If you need to assign, change, or bind multiple IP addresses to an interface, proceed to Step 7.
 - If you require additional VLAN configuration, proceed to Step 8.
 - Otherwise, click **Apply**; the adapter configuration is complete. Proceed to ["Viewing Interface Statistics"](#) on page 1248 for verification.




7. If applicable, assign an IP address, change an IP address, or bind multiple IP addresses to an interface.
 - a. Select the **Physical Interface**.
 - b. Click **Edit**. The Configure Interface IPs dialog displays.
 - c. Click **Add IP**. The Add List Item dialog displays.
 - d. Specify the IP address (IPv4 or IPv6) and subnet mask (for IPv4) or prefix length (for IPv6). An IPv6 address can be link-local or global. Click **OK** to close the dialog.
 - e. Click **OK**.
 - f. Click **Apply**.



8. If applicable, configure Virtual LAN (VLAN) options (see ["About Link Settings"](#) on page 1232):
 - a. By default, the native VLAN ID for any ProxySG appliance interface is **1**, as most switches by default are configured to have their native VLAN IDs as **1**. Only change the **Native VLAN for Interface** value if the native VLAN ID of the switch or router connected to this interface is a value other than **1**; match that value here.
 - b. To add VLANs other than the native VLAN to the interface, click **New VLAN**. The Add IP Address dialog displays.

c.

**Note:**

9. Configure the VLAN options:
 - a. Specify the **VLAN ID** (VID) number of the VLAN accepted on this interface.
 - b. Click **Add IP** to display the Add List Item dialog.
 - c. Specify the VLAN IP address and subnet mask; click **OK** to close the dialog.
 - d. The receiving packet and browser behavior is the same as for physical interfaces (see [Table 67–1, "How each interception option affects connections."](#) on page 1231) with the exception of **Use physical interface setting**, which applies the same configuration to the VLAN as was set on the physical interface.
 - e. Click **OK** in both dialogs.
10. Click **Apply**.

Improve Resiliency or Create a Bigger Pipe with an Aggregate Interface

(Added in SGOS 6.5.9.8) Multiple physical interfaces may be bundled into one logical multi-gigabit aggregate interface using standard 1 GB or 10 GB physical interfaces. This provides increased throughput and network resiliency. If an

interface which is part of an aggregated link goes down, its traffic will move to the other interfaces within the aggregate interface. When the interface comes back up, the traffic is redistributed across all of the links.

An aggregate interface is created on the fly when the first physical interface is added to it. Settings from the first member are applied to the aggregate link (or group). Consequent members take their common settings from the parent aggregate link. Common settings include:


- MTU size
- Reject inbound
- Allow intercept
- VLAN trunking
- Native VLAN
- Spanning tree
- IPv6 auto-linklocal

Editing VLAN settings will update the settings for that specific VLAN on all member interfaces.

Link aggregation is accomplished using the industry-standard IEEE 802.1AX Link Aggregation standard. Switch support and switch configuration are required. The switch and appliance must be cabled port-to-port.

Configure a New Aggregate Interface:

1. Select the **Configuration > Network > Adapters > Aggregate Interfaces** tab.
2. Click **New Aggregate Interface**. The **New Aggregate Interface** pane displays.



3. Select the **Identifier**; this reference is used on the **Interfaces** displays.
4. Optionally, give the interface a intuitive name in the **Label** field.
5. Click each interface you want to add to the aggregate interface.
 - Only available interfaces are displayed. As an example, a physical interface used in a bridge will not be displayed.
 - Up to 32 interfaces can be added to an aggregate interface.
 - The **Link State** column provides link information such as Enable requested, Disabled, Auto <negotiated speed, duplex>, and so on.
 - The MAC address of the first physical interface assigned to the aggregate interface becomes the MAC address of the aggregate link.
 - Interfaces from different adapters may belong to the same aggregate link.
 - Interfaces with different speeds are allowed in an aggregate interface.
 - When a physical link is added to an aggregate interface, the VLAN configurations are merged. A VLAN on the new physical interface is created on the aggregate, and the existing configuration is copied to the aggregate and all other group members; if a VLAN exists on the aggregate interface, it will be created on the new physical interface. If both a new member and the aggregate interface have a VLAN configuration, all VLAN settings are copied from the aggregate VLAN to the member VLAN, except the IP address.

Note: LACP (Link Aggregation Control Protocol) standby link selection and dynamic key management are not supported.

6. Click **OK**. The **New Aggregate Interface** pane closes.
7. Click **Apply** on the **Aggregate Interfaces** tab. The settings update. The **LACP State** column in the **Aggregate Interfaces** panel provides the LACP status, as follows:
 - **Up:** The member is healthy and operates normally from LACP perspective.
 - **Synchronizing:** Peers are out of sync with the port, or unable to exchange LACP PDUs.
 - **Negotiating:** Exchanging key information with the peer. If it persists, this might indicate that the peer is not in the correct link aggregation, or that other configuration on the switch port differs from the rest of the aggregation.
 - **Suspended:** The port is not being used by LACP. Potential reason is port is in half-duplex.
 - **Disabled:** The physical interface is disabled.
 - **Down:** The physical interface has no link.
 - **Creation Pending:** **Apply** hasn't been clicked; the link hasn't been created yet.
8. To verify the aggregate interface, click the **Interfaces** tab; the identifier will now appear under **Aggregate Group**. On the **Aggregate Interfaces** tab, any applicable VLANs appear under the **VLANs on Interface aggr:x** heading.

Note: The list of IP Address in the **VLANs on Interface aggr:x** panel is cumulative; all IP addresses for the aggregate interface are listed.

The screenshot displays the 'Interfaces' configuration page. At the top, there are tabs for 'Interfaces', 'Aggregate Interfaces', and 'Bridges'. The 'Interfaces' tab is selected, showing 'Interfaces & VLANs'. Under 'Physical Interfaces', a table lists three interfaces: 0:0, 1:0 (WAN), and 1:1 (LAN). All three have a 'Link State' of 'Auto: 1 Gbps FDX' and an 'Aggregate Group' of 'aggr:0'. A blue callout box points to the 'aggr:0' column with the text 'Aggregate Group related VLANs'. Below this table is a button 'Configure interface 1:0'. Under 'VLANs on Interface 1:0', there is a table with columns 'VLAN ID', 'IP Address', and 'Prefix Length (Subnet I'. The 'Physical Interface' row shows a prefix length of 24, and the 'VLAN ID' row shows a prefix length of 64.

Remove a Member Interface

1. On the **Aggregate Interfaces** tab, click **Configure interface aggr:x**.
2. On the **Configure Aggregate Interface aggr:x** window which is displayed, clear each interface you want to remove from the aggregate group.
3. Click **OK**.
4. Click **Apply**.

Note: A removed member will maintain the common settings it inherited from the aggregate interface.

Delete an Aggregate Interface

1. On the **Aggregate Interfaces** tab, click **Delete Aggregate Interface**.
2. On the **Confirm delete?** pop up, click **Yes**.
3. Click **Apply**. After the “success” message displays, you can verify that aggregate interface no longer shows in the **Aggregate Interfaces** list.

Notes

- The VLAN information on the **Aggregate Interfaces** tab is for viewing only. Configure VLANs on the **Interfaces** tab.
- See the **Statistics > Network > Interface History** tab to view statistics on an aggregate interface.
- Disabling the aggregate link interface will disable the physical interface of each member of the aggregate group. Individual physical interfaces within the group can be enabled or disabled.
- Packets for any given connection will always be transmitted over the same physical link, limiting the burst speed for the connection to the capacity of that specific link rather than the sum of all the links.
- Link aggregation and bridging can't be configured on the same physical interface, though they may both be used in a single deployment.

Switch Configuration


Link aggregation should work on any switch configured to use LACP. Consult the documentation from your switch vendor when configuring link aggregation.

Viewing Interface Statistics

As traffic flows to and from the ProxySG appliance, you can review statistics for each interface (including VLAN traffic). This allows you to verify your deployment is optimized. For example, if you notice that traffic flowing through the LAN interface is consistently near capacity, you might consider routing traffic differently or spreading the load to another ProxySG appliance.

To view interface-specific statistics:

In the Management Console, select **Statistics > Network > Interface History**.



1. From the **Duration** drop-down list, select a time frame.
2. Select a data type:

Data Type	Description
Bytes Sent	The number of outgoing bytes sent from this interface or VLAN.
Bytes Received	The number of inbound bytes received on this interface or VLAN.
Packets Sent	The number of outgoing packets sent from this interface or VLAN.
Packets Received	The number of inbound packets received on this interface or VLAN.
Input Errors	The number of input and output errors that occurred on the interface (not applicable on VLANs). This information provides details that Symantec Technical Support uses to troubleshoot issues.
Output Errors	

3. Select an interface to view. If an interface has attached VLANs, the tree expands to display the VLAN(s), which are also selectable.

In the graph area, roll your mouse over data lines to view exact metrics.

See Also

[Chapter 72: "Monitoring the ProxySG" on page 1299](#)

Detecting Network Adapter Faults

The ProxySG appliance can detect whether the network adapters in an appliance are functioning properly. If the appliance detects a faulty adapter, it stops using it. When the fault is remedied, the ProxySG appliance detects the functioning adapter and uses it normally.

To determine whether an adapter is functioning properly:

1. Check whether the link is active (that is, a cable is connected and both sides are up).
2. Check the ratio of error packets to good packets: both sent and received.
3. Check if packets have been sent without any packets received.
4. Check the event log. If an adapter fault is detected, the ProxySG appliance logs a severe event. In addition, the ProxySG appliance logs an entry even when a faulty adapter is restored.

Chapter 68: Software and Hardware Bridges

This section describes the SGOS hardware and software bridging capabilities. Network bridging through the ProxySG provides transparent proxy pass-through and failover support.

Topics in this Section

This section contains the following topics:

- ❑ ["About Bridging"](#)
- ❑ ["About the Pass-Through Adapter" on page 1254](#)
- ❑ ["Configuring a Software Bridge" on page 1255](#)
- ❑ ["Configuring Programmable Pass-Through/NIC Adapters" on page 1256](#)
- ❑ ["Customizing the Interface Settings" on page 1258](#)
- ❑ ["Setting Bandwidth Management for Bridging" on page 1258](#)
- ❑ ["Configuring Failover" on page 1259](#)
- ❑ ["Bridging Loop Detection" on page 1261](#)
- ❑ ["Adding Static Forwarding Table Entries" on page 1263](#)
- ❑ ["Bypass List Behavior" on page 1264](#)

About Bridging

A bridge is a network device that interconnects multiple computer networks. Unlike a hub, a bridge uses the Ethernet frame's destination MAC address to make delivery decisions. Because these decisions are based on MAC addressing, bridges are known as Layer 2 devices. This Layer 2 functionality is similar to that used by switches. Bridging is especially useful in smaller deployments in which explicit proxies or L4 switches are not feasible options.

Bridging functionality allows each ProxySG to be easily deployed as a transparent redirection device, without requiring the additional expense and maintenance of L4 switches or WCCP-capable routers. Transparent bridges are deployed in-path between clients and routers—all packets must pass through them, though clients are unaware of their presence.

A branch office that would take advantage of a bridging configuration is likely to be small; for example, it might have only one router and one firewall in the network, as shown below.




Figure 68–1 A Bridged Configuration

To ensure redundancy, the ProxySG supports both serial and parallel failover modes. See ["Configuring Failover"](#) on page 1259 for more information about serial and parallel failover configurations.

About Bridging Methods

The ProxySG provides bridging functionality by two methods:

- ❑ Software—A software, or *dynamic*, bridge is constructed using a set of installed interfaces. Within each logical bridge, interfaces can be assigned or removed. In the event of failure, software bridges fail closed—traffic is not passed. This behavior can be desirable if you want to pass traffic to a redundant ProxySG and/or link.

See ["Configuring Programmable Pass-Through/NIC Adapters"](#) on page 1256 for more information.

- ❑ Hardware—A hardware, or *pass-through*, bridge uses a dual interface Ethernet adapter. This type of bridge provides pass-through support—in the event of failure, traffic passes through the appliance.

See ["About the Pass-Through Adapter"](#) on page 1254 for more information.

Note: If you want to use an L4 switch or an explicit proxy instead of bridging, you must disable the pass-through card.

Traffic Handling

Bridges are used to segment Ethernet collision domains, thus reducing frame collisions. To make efficient delivery decisions, the bridge must discover the identity of systems on each collision domain. The bridge uses the source MAC address of frames to determine the interface that the device can be reached from and stores that information in the bridge forwarding table. When packets are received, the bridge consults the forwarding table to determine which interface to deliver the packet to. The only way to bypass the bridge forwarding table lookup is to define a static forwarding entry. For more information on static forwarding entries, see ["Adding Static Forwarding Table Entries"](#) on page 1263.

Trust Destination MAC

When the ProxySG is in transparent bridging mode, the ProxySG always “trusts” the destination MAC address of inbound packets and does not consult its routing table. Trust Destination MAC is enabled by default (when the ProxySG is in transparent bridging mode) and cannot be disabled. For more information on Trust Destination MAC, see ["Routing on the ProxySG"](#) on page 795.

About Bridging and Policy

Because the bridge intercepts all traffic, you can take advantage of the powerful proxy services and policies built into the ProxySG to control how that traffic is handled. If the ProxySG recognizes the intercepted traffic, you can apply policy to it. Unrecognized traffic is forwarded out. The following diagram illustrates this traffic handling flow.




Figure 68–2 Traffic Flow Decision Tree

Because policy can be applied only to recognized protocols, it is important to specify port ranges that will capture all traffic, even that operating on lesser-known ports.

About the Pass-Through Adapter

A pass-through adapter is a dual interface Ethernet adapter designed by ProxySG to provide an efficient fault-tolerant bridging solution. If this adapter is installed on a ProxySG, SGOS detects the adapter on system bootup and automatically creates a bridge—the two Ethernet interfaces serve as the bridge ports. If the ProxySG is powered down or loses power for any reason, the bridge fails open; that is, network traffic passes from one Ethernet interface to the other. Therefore, network traffic is uninterrupted, but does not route through the appliance.

Important: This scenario creates a security vulnerability.

After power is restored to the ProxySG, the bridge comes back online and network traffic is routed to the appliance and thus is subject to that appliance's configured features, policies, content scanning, and redirection instructions. Bridging supports only failover; it does not support load balancing.

Note: The adapter state is displayed on **Configuration > Network > Adapters**.

Deployment Recommendations

ProxySG recommends racking and cabling the ProxySG while it is powered off. This enables you to confirm that the pass-through adapter is functioning and that traffic is passing through the appliance. If traffic is not being passed, confirm that you have used the correct cabling (crossover or straight).

Reflecting Link Errors

When the ProxySG is deployed transparently with bridging enabled, link errors that occur on one interface can be reflected to the other bridge interface. This allows a router connected to the ProxySG on the healthy link to detect this failure and recompute a path around this failed segment. When the interface with the original link error is brought back up, the other interface is automatically restarted as part of the health check process.

Reflecting link errors requires that two interfaces be available and connected in a bridging configuration; it also requires that the `propagation-failure` option is enabled. By default, `propagation-failure` is disabled.

Note: This feature is only applicable to a two-interface hardware or software bridge. The `propagation-failure` option sets itself to disabled in any other scenario.

If the link goes down while `propagation-failure` is disabled, the previous link state is immediately reflected to the other interface if `propagation-failure` is enabled during this time.


Configuring a Software Bridge

This section describes how to link adapters and interfaces to create a network bridge.


Before configuring a software bridge, ensure that your adapters are of the same type and use the same settings. Although the software does not restrict you from configuring bridges with adapters of different speeds and MTU configurations (for example, ports speeds of 10/100 Mbit/s and 1 GigE combined with an MTU of 1400 and 1500, respectively), the resulting behavior is unpredictable.

To create and configure a software bridge:

1. Select **Configuration > Network > Adapters > Bridges**.
2. Click **New**. The Create Bridge dialog displays.



3. Configure bridge options:
 - a. In the **Bridge Name** field, enter a name for the bridge—up to 16 characters. The bridge name is case insensitive, that is, you cannot name one bridge **ABC** and another bridge **abc**.
 - b. (Optional) If you want to assign the bridge to a failover group select it from the **Failover Group** drop-down list.
 - c. See "[Configuring Failover](#)" on page 1259 for more information about configuring failover.
 - d. Click **Add**. The **Add Bridge Interface** dialog displays.



4. Configure the bridge interface options:
 - a. From the **Interface** drop-down menu, select an interface.
 - b. (Optional) To enable bridging loop avoidance, select **Enable Spanning Tree**. See ["Bridging Loop Detection"](#) on page 1261 for more information about the Spanning Tree Protocol.
 - c. If you are using firewall configurations that require the use of static forwarding table entries, add a static forwarding table entry that defines the next hop gateway that is on the correct side of the bridge. For more information on static forwarding table entries, see ["Adding Static Forwarding Table Entries"](#) on page 1263.
 - d. Click **OK**.
 - e. Repeat Step 4 for each interface you want to attach to the bridge.
5. Click **OK** to close the Create Bridge Interface and Create Bridge dialogs.
6. Click **Apply**.

Configuring Programmable Pass-Through/NIC Adapters

Some ProxySG appliances ship (when ordered) with a network adapter card that can be used as a pass-through adapter or as a Network Interface Card (NIC), depending on the configured mode. If the network adapter mode is set to disabled, the adapter interfaces can be used as NICs or as part of a software bridge.

If your appliance includes a programmable adapter card, the following programmable adapter modes are available:


- ❑ **Disabled**—Disables the bridge and allows the adapter interfaces to be reused as NICs or as part of another bridge.
- ❑ **Fail Open**—If the ProxySG fails, all traffic passes through the bridge so clients can still receive data.
- ❑ **Fail Closed**—If the ProxySG fails, all traffic is blocked and service is interrupted. This mode provides the same functionality as a user-configured software bridge.

Note: If you create a software bridge, the programmable bridge card mode is implicitly **Fail Closed** (if the appliance fails, the software bridge is non-functional).

The following procedure describes programmable adapter configuration.

To configure the function of the programmable adapter:

1. Select **Configuration > Network > Adapters > Bridges**.
2. In the **Bridges** section, select the bridge you want to configure.
3. Click **Edit**. The Edit Bridge dialog displays.



4. Configure the bridge options:
 - a. Select the desired mode from the **Mode** drop-down list.

- b. If you have a two-interface bridge and want to enable link error propagation, select the **Propagate Failure** check box.
 - c. (Optional) Click **Clear Bridge Statistics** to reset the traffic history of the bridge, which includes packet and byte counts, to 0.
 - d. Click **OK** to save your changes and close the Edit Bridge dialog.
5. Click **Apply**.

Customizing the Interface Settings

To further customize the bridge, edit the interface settings.

Editing the interface settings allows you to

- ❑ Allow transparent interception. It is bypassed by default. You must configure the WAN interface to allow transparent interception.

Note: If you have a MACH5 license, a programmable bridge card, and labeled WAN/LAN interfaces, the WAN interface allows transparent interception by default.

- ❑ Firewall incoming traffic. Firewalls must be specifically configured.

See "[Configuring Adapters and Virtual LANs](#)" on page 1229 for more information.


The **Bridge Settings** options allow you to clear bridge forwarding table and clear bridge statistics.

Setting Bandwidth Management for Bridging

After you have created and configured a bandwidth management class for bridging (**Configuration > Bandwidth Mgmt. > BWM Classes**), you can manage the bandwidth used by all bridges. See "[Configuring Bandwidth Allocation](#)" on page 602 for more information on bandwidth management.

To configure bandwidth management for bridging:

1. Select **Configuration > Network > Adapters > Bridges**.



2. In the **Bridging Bandwidth Class** drop-down menu, select a bandwidth management class to manage the bandwidth for bridging, or select **<none>** to disable bandwidth management for bridging.

Note: This setting only controls the bandwidth class used by bypassed traffic on this bridge. To manage intercepted traffic, you must define a Manage Bandwidth policy (using VPM or CPL).

3. Click **Apply**.

Configuring Failover

In failover mode, two appliances are deployed, a master and a slave. The master sends keepalive messages (*advertisements*) to the slave appliance. If the slave does not receive advertisements at the specified interval, the slave takes over for the master. When the master comes back online, the master takes over from the slave again.


The SGOS bridging feature allows two different types of failover modes, *parallel* and *serial*. Hardware and software bridges allow different failover modes:

- ❑ Software bridges allow serial or parallel failover. However, note that if the ProxySG fails, serial failover also fails.
- ❑ Hardware bridges allow serial or parallel failover.

Parallel Failover

In parallel failover mode, two systems are deployed side by side on redundant paths. In parallel failover, the slave does not actively bridge any packets unless the master fails. If the master fails, the slave takes over the master IP address and begins bridging. A parallel failover configuration is shown in the following figure.


Because of the redundant paths, you must enable Spanning Tree to avoid bridge loops. See "[Bridging Loop Detection](#)" on page 1261 for more information about STP.



Serial Failover

In serial failover mode, the slave is in-path and continuously bridges packets, but does not perform any other operations to the bridged traffic unless the master fails. If the master fails, the slave takes over the master IP address and applies policy, etc. A serial configuration is shown in the following figure.

If you are relying on a hardware bridge for serial failover, you must configure the pass-through bridge to be in fail open mode. See ["Configuring Programmable Pass-Through/NIC Adapters"](#) on page 1256 for more information about configuring bridge modes.



Configuring Failover

Failover is accomplished by doing the following:

- ❑ Creating virtual IP addresses on each proxy.
- ❑ Creating a failover group.
- ❑ Attach the failover group to the bridge configuration.
- ❑ Selecting a failover mode (parallel or serial - this can only be selected using the CLI).

Both proxies can have the same priority (for example, the default priority). In that case, priority is determined by the local IP address—the ProxySG with the highest local IP will assume the role of master.

Example

The following example creates a bridging configuration with one bridge on standby.

Note: This deployment requires a hub on both sides of the bridge or a switch capable of interface mirroring.

- ❑ ProxySG A—software bridge IP address: 10.0.0.2. Create a virtual IP address and a failover group, and designate this group the *master*.

```
SGOS_A#(config) virtual-ip address 10.0.0.4
SGOS_A#(config) failover
SGOS_A#(config failover) create 10.0.0.4
SGOS_A#(config failover) edit 10.0.0.4
SGOS_A#(config failover 10.0.0.4) master
SGOS_A#(config failover 10.0.0.4) enable
```

The preceding commands create a failover group called 10.0.0.4. The priority is automatically set to 254 and the failover interval is set to 40.

- ❑ ProxySG B—software bridge IP address: 10.0.0.3. Create a virtual IP address and a failover group.

```
SGOS_B#(config) virtual-ip address 10.0.0.4  
SGOS_B#(config) failover  
SGOS_B#(config failover) create 10.0.0.4  
SGOS_B#(config failover) edit 10.0.0.4  
SGOS_B#(config failover 10.0.0.4) enable
```

In the bridge configuration on each SG, attach the bridge configuration to the failover group:

```
SGOS_A#(config bridge bridge_name) failover group 10.0.0.4  
SGOS_B#(config bridge bridge_name) failover group 10.0.0.4
```

□ Specify the failover mode:

```
SGOS_A#(config bridge bridge_name) failover mode serial  
SGOS_B#(config bridge bridge_name) failover mode serial
```

Bridging Loop Detection


Bridging now supports the Spanning Tree Protocol (STP). STP is a link management protocol that prevents bridge loops in a network that has redundant paths that can cause packets to be bridged infinitely without ever being removed from the network.

STP ensures that a bridge, when faced with multiple paths, uses a path that is loop-free. If that path fails, the algorithm recalculates the network and finds another loop-free path.


The administrator can enable or disable spanning tree participation for the interface.

Enable spanning tree participation:

1. Select **Configuration > Network > Adapters > Bridges**.
2. Select the desired bridge.
3. Click **Edit**. The Edit Bridge dialog displays.



4. Select the interface to configure and click **Edit**. The Edit Bridge Interface dialog displays.



5. Select **Enable Spanning Tree**.
6. Click **OK** to close the Edit Bridge Interface and Edit Bridge dialogs.
7. Click **Apply**.


Adding Static Forwarding Table Entries

Certain firewall configurations require the use of static forwarding table entries. These firewall failover configurations use virtual IP (VIP) addresses and virtual MAC (VMAC) addresses. When a client sends an ARP request to the firewall VIP, the firewall replies with a VMAC (which can be an Ethernet multicast address); however, when the firewall sends a packet, it uses a physical MAC address, not the VMAC.

The solution is to create a static forwarding table entry that defines the next hop gateway that is on the correct side of the bridge.

To create a static forwarding table:

1. Select **Configuration > Network > Adapters > Bridges**.
2. Select the bridge to edit and click **Edit**. The Edit Bridge Interface dialog displays.



3. Add the static forwarding table entry.
 - a. In the Edit Bridge dialog, select the interface on which to create the static forwarding table entry.
 - b. Click **Edit**.
 - c. In the Edit Bridge Interfaces dialog, click **Add**.
 - d. In the Add MAC dialog, add the MAC address of the next hop gateway and click **OK**.
4. Click **OK** to close the Edit Bridge Interface and Edit Bridge dialogs.
5. Click **Apply**.

Bypass List Behavior

Starting with SGOS 5.x, static and dynamic bypass operates differently, depending on how the ProxySG intercepts the traffic, as follows:

- ❑ When the ProxySG is installed in a bridging deployment, bridging is used for bypass.
- ❑ When the ProxySG is installed as a router or external layer 4 load balancers are used to redirect traffic to the ProxySG, routing is used for bypass, but only if IP Forwarding is enabled.

Otherwise, traffic is dropped instead of being bypassed.

- ❑ When the ProxySG is installed in a WCCP deployment, either Generic Route Encapsulation (GRE) or Layer 2 (L2) redirection is used for bypass.

To understand this process, review the following information:

- *SGOS versions before 5.4 only.* If L2 redirection was used in earlier SGOS releases to forward packets from the router to the ProxySG, the ProxySG did not always treat those packets as arriving by WCCP, so static and dynamic bypass never attempted to use WCCP packet return.

In those configurations, IP Forwarding had to be enabled so packets were properly returned to the router. Otherwise, the traffic was dropped.

- SGOS 5.4 and later overcomes this limitation and properly uses WCCP packet return to redirect bypassed traffic back to the router, supporting the following combination of packet forwarding and return options:

Packet forwarding	Packet return
GRE	GRE
L2	GRE
L2	L2

To set these options in the Management Console, select **Configuration > Network > WCCP**. Select **Enable WCCP** and click **New**. Enter the required prerequisite information (such as Service Group, Priority, and so on) and select options for **Forwarding Type** and **Returning Type**.

For additional details, click **Cancel** in the New Service dialog and click **Help** on the **WCCP** tab. The corresponding CLI commands are discussed in *Command Line Interface Reference*.

For more details, consult the *Release Notes* for SGOS version.

Chapter 69: Configuring Management Services

This section describes how to configure administrative access to the ProxySG consoles, including the Management Console and the command line interface (CLI). It includes the following topics:

- ❑ ["Overview of Management Services"](#) on page 1267
- ❑ ["Creating a Management Service"](#) on page 1268
- ❑ ["Managing the HTTP Console"](#) on page 1269
- ❑ ["Managing the HTTPS Console \(Secure Console\)"](#) on page 1270
- ❑ ["Managing the SNMP Console"](#) on page 1273
- ❑ ["Managing the SSH Console"](#) on page 1273
- ❑ ["Managing the Telnet Console"](#) on page 1277

Overview of Management Services

The ProxySG provides administrative access to the appliance through management services, or *consoles*. The following management services are available:

- ❑ HTTP and HTTPS Consoles: These consoles are designed to allow you access to the Management Console. The HTTPS Console is created and enabled by default; the HTTP Console is created by default but not enabled because it is less secure than HTTPS.
- ❑ SSH Console: This console is created and enabled by default, allowing you access to the CLI using an SSH client.
- ❑ SNMP Console: This console is created by default, but disabled. SNMP listeners set up the UDP and TCP ports the ProxySG uses to listen for SNMP commands.
- ❑ Telnet Console: This console not created by default because the passwords are sent unencrypted from the client to the ProxySG, which is less secure than the other management services. You must create and enable the Telnet console service before you can access the appliance through a Telnet client (not recommended).

Table 69–1 Management Services

Management Service	Default Port	Status	Configuration Discussed
HTTPS-Console	8082	Enabled	"Managing the HTTPS Console (Secure Console)" on page 1270.
SSH-Console	22	Enabled	"Managing the SSH Console" on page 1273
HTTP-Console	8081	Disabled	"Managing the HTTP Console" on page 1269

Table 69–1 Management Services (Continued)


Management Service	Default Port	Status	Configuration Discussed
SNMP	161	Disabled	"Managing the SNMP Console" on page 1273
Telnet-Console	—	Not Created	"Managing the Telnet Console" on page 1277

Creating a Management Service

Management services are used to manage the ProxySG. As such, bypass entries are ignored for connections to console services. For more information, see ["Overview of Management Services"](#) on page 1267.

To edit or create a management service:


1. Select the **Configuration > Services > Management Services** tab.



Name	Service	Proxy IP	Port	Enabled
HTTP-Console	HTTP Console	<All>	8081	☒
HTTPS-Console	HTTPS Console	<All>	8082	☒
SSH-Console	SSH Console	<All>	22	☒
SNMP	SNMP	<All>	161	☐
Telnet-Console	Telnet Console			☐

2. To enable or disable a service, select or de-select the **Enable** option.
3. To change other settings on a specific console, highlight the service and click **Edit**.
4. To create a new console service, click **New**.

Note: The HTTP Console is used in this example.



5. Enter a meaningful name in the **Name** field.
6. From the **Console** drop-down list, select the console that is used for this service.
7. Configure the new listener options:
 - a. Click **New** to view the **New Listener** dialog. A listener defines the fields where the console service will listen for traffic.
 - b. Select a destination option:
 - **All ProxySG IP addresses**—indicates that service listens on all addresses (IPv4 and IPv6).
 - **IP Address**—indicates that only destination addresses match the IP address. IPv4 or IPv6 addresses can be specified. Note that when IPv6 addresses are specified, they must be global (not linklocal).
 - c. **Port**—Identifies the port you want this service to listen on. Port 8081 is the default port.
 - d. **Enabled**—Select this option to enable the listener.
 - e. Click **OK** to close the New Listener dialog.
8. Click **OK** to close the New Service dialog.
9. Click **Apply**.

Managing the HTTP Console

The default HTTP Console is already configured; you only need to enable it.

You can create and use more than one HTTP Console as long as the IP address and the port are unique.

Administrative access to the appliance for the HTTP console can be controlled with the following authentication types:

- ❑ The predefined **admin** account
- ❑ Local authentication realm
- ❑ Certificate authentication realm
- ❑ IWA authentication realm (with basic authentication, secured with TLS)
- ❑ LDAP authentication realm (secured with TLS)

To create a new HTTP Console service or edit an existing one, see "[Creating a Management Service](#)" on page 1268.

Managing the HTTPS Console (Secure Console)

The HTTPS Console provides secure access to the Management Console through the HTTPS protocol.

You can create multiple management HTTPS consoles, allowing you to simultaneously access the Management Console using any IP address belonging to the ProxySG as well as any of the appliance's virtual IP (VIP) addresses. The default is HTTPS over port 8082.

Administrative access to the appliance for the HTTPS console can be controlled with the following authentication types:

- ❑ The predefined **admin** account
- ❑ Local authentication realm
- ❑ Certificate authentication realm (refer to the *Common Access Card Solutions Guide* for information)
- ❑ IWA authentication realm (with basic authentication, secured with TLS)
- ❑ LDAP authentication realm (secured with TLS)

Creating a new HTTPS Console service requires three steps, discussed in the following sections:

- ❑ Selecting a keyring (a key pair and a certificate that are stored together)
- ❑ Selecting an IP address and port on the system that the service will use, including virtual IP addresses
- ❑ Enabling the HTTPS Console Service

Selecting a Keyring

The ProxySG ships with a default keyring that can be reused with each secure console that you create. You can also create your own keyrings.

To use the default keyring, accept the default keyring through the Management Console. If using the CLI, the default keyring is automatically used for each new HTTPS Console that is created. To use a different keyring you must edit the console service and select a new keyring using the attribute keyring command.

Note: If you get “host mismatch” errors or if the security certificate is called out as invalid, create a different certificate and use it for the HTTPS Console. For more information on keyrings and certificates, see [Chapter 61: "Managing X.509 Certificates"](#) on page 1113.

For information on creating a key pair and a certificate to make a keyring, see [Chapter 61: "Managing X.509 Certificates"](#) on page 1113.

Selecting an IP Address


You can use any IPv4 or IPv6 address on the ProxySG for the HTTPS Console service, including virtual IP addresses. Note that when IPv6 addresses are specified, they must be global (not linklocal). For information on how to create a virtual IP address, see ["Creating a VIP"](#) on page 823.

Enabling the HTTPS Console Service

The final step in editing or creating an HTTPS Console service is to select a port and enable the service.

To create or edit an HTTPS Console port service:

1. Select the **Configuration > Services > Management Services** tab.
2. Perform one of the following:
 - To create a new HTTPS Console service, see ["Creating a Management Service"](#) on page 1268.
 - To edit the configuration of an existing HTTPS Console service, highlight the HTTPS Console and click **Edit**. The Edit Service dialog displays.



3. From the **Keyring** drop-down list, which displays a list of existing keyrings on the system, select a keyring. The system ships with a default keyring that is reusable for each HTTPS service.

Note: You cannot use the configuration-passwords-key keyring or the application-key keyring for console services. In addition, you should remove unwanted cipher suites from the keyring used to make SSL connections. See ["Editing or Creating an SSL Device Profile"](#) on page 1170.

4. Select SSL/TLS protocols:
 - a. Select **TLSv1.2**, **TLSv1.1**, and **TLSv1**.
 - b. If **SSL v3** or **SSL v2** are selected, clear the selection. TLS versions 1, 1.1, and 1.2 have superseded SSL Versions 2 and 3.
5. (If configuring CAC authentication or a certificate realm) Select **Verify Client**. This setting enables mutual SSL authentication for the Management Console. For more information about mutual SSL authentication, see ["About Mutual SSL Authentication"](#) on page 325.
6. Configure the new listener options:
 - a. Click **New** to view the **New Listener** dialog. A listener defines the fields where the console service will listen for traffic.
 - b. Select a destination option:
 - **All ProxySG IP addresses**—Indicates that service listens on all addresses (IPv4 and IPv6).

- **IP Address**—Indicates that only destination addresses match the IP address. You can enter an IPv4 or an IPv6 address. Note that when IPv6 addresses are specified, they must be global (not linklocal).
- c. **Port**—Identifies the port you want this service to listen on. Port 8081 is the default port.
- d. **Enabled**—Select this option to enable the listener.
- e. Click **OK** to close the New Listener dialog.
- 7. Click **OK** to close the Edit Service dialog.
- 8. Click **Apply**.

Creating a Notice and Consent Banner for the Management Console

You can install Content Policy Language (CPL) to create a Notice and Consent banner for the Management Console.

Refer to the *Notice and Consent Banner Configuration Webguide* for more information.

Managing the SNMP Console

There is one disabled SNMP listener defined by default on the ProxySG, which you can delete or enable, as needed. You can also add additional SNMP services and listeners. Enabling SNMP listeners sets up the ProxySG IPv4/IPv6 addresses and ports (UDP and TCP) on which the ProxySG listens for SNMP commands. The SNMP console supports passphrase authentication. Other authentication types are not supported.

To create and enable an SNMP service:

1. Select the **Configuration > Services > Management Services** tab.
2. Click **New**. The New Service dialog displays.
3. Follow steps 2–5 in the section titled "[Creating a Management Service](#)" on page 1268.

Managing the SSH Console

By default, the ProxySG uses Secure Shell (SSH) and password authentication so administrators can access the CLI securely. SSH is a protocol for secure remote logon over an insecure network.

Authentication for administrative users connecting to the appliance via SSH can be controlled with the following authentication types:

- ☐ The predefined **admin** account
- ☐ Local authentication realm
- ☐ IWA authentication realm
- ☐ LDAP authentication realm
- ☐ Localized RSA key

When managing the SSH console, you can:

- ❑ Enable or disable a version of SSH
- ❑ Generate or re-generate SSH host keys
- ❑ Create or remove client keys and director keys
- ❑ Specify a welcome message for clients accessing the ProxySG using SSHv2.

To create a new SSH Console service or edit an existing one, see "[Creating a Management Service](#)" on page 1268.

Managing the SSH Host Key Pairs

You can manage the SSH host connection either through the Management Console or the CLI.

Note: By default, SSHv2 is enabled and assigned to port 22. You do not need to create a new host key unless you want to change the existing configuration. SSHv1 is disabled by default.

To manage the SSH host:

1. Select the **Configuration > Authentication > Console Access > SSH Host** tab.

The screenshot shows the 'SSH Host' configuration page. It includes tabs for 'Console Account', 'Console Access', 'SSH Host' (active), and 'SSH Client'. Below the tabs, there are three main sections:

- SSHv1 Host Key Pair:** Contains a 'Create' button, a 'Delete' button, and a text field with the value '1024 35 141473172771263863772574660229091912141339481229'.
- SSHv2 Host Key Pair:** Contains a 'Create' button, a 'Delete' button, and a text field with the value 'ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAqM7yY4FsaK6eAT4fG5f0'.
- SSHv2 Welcome Banner:** Contains a text field with the value 'You are accessing ProxySG-22, located in the Tower 2 Server Room.'

To delete a host key pair:

Click the **Delete** button for the appropriate version of SSH.

The key pair is deleted and that version of SSH is disabled.

Note: If you disable both SSHv1 and SSHv2, you could be locked out of the CLI, requiring you to re-create an SSH key pair using the terminal console. (You can re-create the SSH keys through the Management Console.)

```
SGOS (config ssh-console) create host-keypair {sshv1 | sshv2 | <Enter>}
```

To create a host key pair:

Click the **Create** button for the appropriate version of SSH.

The new key pair is created and that version of SSH is enabled. The new key pair is displayed in the appropriate pane.

Note: If you receive an error message when attempting to log in to the system after regenerating the host key pair, locate the `ssh known hosts` file and delete the system's IP address entry.

Creating a Notice and Consent Banner for SSH

You can create a Notice and Consent banner for the SSH console using the `sshv2-welcome-banner` CLI command.

Refer to the *Notice and Consent Banner Configuration Webguide* for more information.

Managing SSH Client Keys

You can import multiple RSA client keys on the ProxySG to provide public key authentication, an alternative to using password authentication. An RSA client key can only be created by an SSH client and then imported onto the ProxySG. Many SSH clients are commercially available for UNIX and Windows.

After you create an RSA client key following the instructions of your SSH client, you can import the key onto the ProxySG using either the Management Console or the CLI. (For information on importing an RSA key, see ["To import RSA client keys:"](#) on page 1276.)

For more information, see one of the following sections:

- ❑ ["About the OpenSSH.pub Format"](#)
- ❑ ["Importing RSA Client Keys"](#) on page 1276

About the OpenSSH.pub Format

Blue Coat supports the `OpenSSH.pub` format. Keys created in other formats will not work.

An `OpenSSH.pub` public key is similar to the following:

```
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAwFI78MKyvL8DrFgcVxpNRHMFkJrBMeBn
2PKcv5oAJ2qz+uZ7hiv7Zn43A6hXwY+DekhtNLOk3HCWmgsrDBE/NOOEnDpLQjBC6t/
T3cSQKZjh3NmBbpE4U49rPduiufvWkuoEiHUb5ylzRGdXRSNJHxxmg5LiGEiKaoELJfsD
Mc= user@machine
```

The `OpenSSH.pub` format appends a space and a user ID to the end of the client key.

The user ID used for each key must be unique.

Notes:

- ❑ 4096 bits is the maximum supported key size.
- ❑ An `ssh-rsa` prefix must be present.
- ❑ Trailing newline characters must be removed from the key before it is imported.

Importing RSA Client Keys


This section discusses how to import RSA client keys into the Management Console to provide more secure authentication compared to user name/password authentication. For more information, see ["Managing SSH Client Keys"](#) on page 1275.

To import RSA client keys:

1. From your SSH client, create a client key and copy it to the clipboard.

Note: The ProxySG cannot create client keys. You must use your SSH client to create a key.

2. Select the **Configuration > Authentication > Console Access > SSH Client** tab.



3. Click **Import**. The Import Client Key dialog displays.
4. Associate a user with a client key:
 - a. Specify whether the client key is associated with an existing user or a new user, and enter the name.

- b. Paste the RSA key that you previously created with an SSH client into the **Client key** field. Ensure that a key ID is included at the end. Otherwise, the import fails.
- c. Click **OK**.

In the **SSH Client** tab, the fingerprint (a unique ID) of the imported key displays.

The screenshot shows a web interface with a tabbed menu at the top: 'Console Account', 'Console Access', 'SSH Host', and 'SSH Client'. The 'SSH Client' tab is selected. Below the tabs, there are two main sections: 'Import Client Key' and 'Delete Client Key'. The 'Import Client Key' section contains an 'Import' button and a text field labeled 'New SSH Client Key'. The 'Delete Client Key' section contains three text fields: 'Username' with the value 'director', 'Keypair' with the value 'director@200-2234567.00d0.8303.3c47.sshv2', and 'Fingerprint' with the value '8A:8F:3D:54:61:64:6B:FA:10:EF:80:43:7A:E0:54:37'.

5. Click **Apply**.

Managing the Telnet Console

The Telnet console allows you to connect to and manage the ProxySG using the Telnet protocol. Remember that Telnet is a clear text protocol that provides no integrity protection. Using Telnet for administrative access will result in administrative credentials being sent in clear text. By default, the Telnet Console is not created.

Authentication for administrative users connecting to the appliance via Telnet can be controlled with the following authentication types:

- ☐ The predefined **admin** account
- ☐ Local authentication realm
- ☐ IWA authentication realm
- ☐ LDAP authentication realm

Blue Coat Systems recommends against using Telnet because of the security hole it creates.

Note: If you enable the Telnet console, be aware that you cannot use Telnet to access all options available in the CLI. Some modules, such as SSL, respond with the error message:

```
Telnet sessions are not allowed access to ssl commands.
```

By default a Telnet shell proxy service exists on the default Telnet port (23). Since only one service can use a specific port, you must delete the shell service if you want to create a Telnet console. Be sure to apply any changes before continuing. If you want a Telnet shell proxy service in addition to the Telnet console, you can re-create it later on a different port. For information on the Telnet service, see [Chapter 15: "Managing Shell Proxies"](#) on page 313.

To create a new Telnet console service or edit an existing one, see "[Creating a Management Service](#)" on page 1268.

Note: To use the Telnet shell proxy (to communicate with off-proxy systems) *and* retain the Telnet Console, you must either change the Telnet shell proxy to use a transparent Destination IP address, or change the destination port on either the Telnet Console or Telnet shell proxy. Only one service is permitted on a port. For more information on the Telnet shell proxy, see [Chapter 15: "Managing Shell Proxies"](#) on page 313.

Chapter 70: Preventing Denial of Service Attacks

This section describes how the ProxySG prevents attacks designed to prevent Web services to users.

Topics in this Section

This section includes the following topics:

- ❑ ["About Attack Detection"](#)
- ❑ ["Configuring Attack-Detection Mode for the Client"](#) on page 1280
- ❑ ["Configuring Attack-Detection Mode for a Server or Server Group"](#) on page 1288

About Attack Detection

The ProxySG appliance can reduce the effects of denial of service (DoS) and distributed-DoS (DDoS) attacks.

DoS and DDoS attacks occur when one or more machines coordinate an attack on a specific Web site in order to cripple or disrupt host services. As the attack progresses, the target host shows decreased responsiveness and often stops responding. Legitimate HTTP traffic is unable to proceed because the infected system no longer has the resources to process new requests.

The ProxySG prevents attacks by limiting the number of simultaneous TCP connections and/or excessive repeated requests from each client IP address that can be established within a specified time frame. If these limits are met, the ProxySG either does not respond to connection attempts from a client already at this limit or resets the connection. It can also be configured to limit the number of active connections to prevent server overloading.

If the ProxySG starts seeing a large number of failed requests, and that number exceeds the configured error limit, subsequent requests are blocked and the proxy returns a warning page.

Failed requests, by default, include various HTTP response failures such as 4xx client errors (excluding 401 and 407) and 5xx server errors. The HTTP responses that you want treated as failures can be so defined by creating policy.

If the requests continue despite the warnings, and the rate exceeds the warning limits that have been specified for the client, the client is then blocked at the TCP level.

You can configure attack detection for both clients and servers or server groups, such as <http://www.bluecoat.com>. The *client* attack-detection configuration is used to control the behavior of attacking sources. The *server* attack-detection configuration is used when an administrator wants to prevent a server from becoming overloaded by limiting the number of outstanding requests that are allowed.

This feature is only available through the CLI. You cannot use the Management Console to enable attack detection.

Configuring Attack-Detection Mode for the Client

To enter attack-detection mode for the client:

From the (config) prompt, enter the following commands:

```
SGOS#(config) attack-detection  
SGOS#(config attack-detection) client
```

The prompt changes to:

```
SGOS#(config client)
```

Changing Global Settings

The following defaults are global settings, used if a client does not have specific limits set. They do not need to be changed for each IP address/subnet if they already suit your environment:

- ☐ client limits enabled: false
- ☐ client interval: 20 minutes
- ☐ block-action: drop (for each client)
- ☐ concurrent-request-limit: unlimited (for each client)
- ☐ connection-limit: 100 (for each client)
- ☐ failure-limit: 50 (for each client)
- ☐ monitor-only: disabled
- ☐ request-limit: unlimited (for each client)
- ☐ unblock-time: unlimited (for each client)
- ☐ warning-limit: 10 (for each client)

To change the global defaults:

Remember that enable/disable limits and interval affect all clients barring instances where limits are enabled and configured for individual clients.

Note: If you edit an existing client's limits to a smaller value, the new value only applies to new connections to that client. For example, if the old value was 10 simultaneous connections and the new value is 5, existing connections above 5 are not dropped.

```

SGOS#(config client) enable-limits | disable-limits
SGOS#(config client) interval minutes
SGOS#(config client) block ip_address [minutes] | unblock ip_address
SGOS#(config client) default block-action drop | send-tcp-rst
SGOS#(config client) default connection-limit
integer_between_1_and_65534
SGOS#(config client) default concurrent-request-limit
integer_between_1_and_2147483647
SGOS#(config client) default failure-limit integer_between_1_and_500
SGOS#(config client) default monitor-only
SGOS#(config client) no default monitor-only
SGOS#(config client) default request-limit
integer_between_1_and_2147483647
SGOS#(config client) default unblock-time minutes_between_1_and_1440
SGOS#(config client) default warning-limit integer_between_1_and_100

```

Table 70–1 Changing Global Defaults

enable-limits disable-limits		Toggles between true (enabled) and false (disabled). The default is false. This is a global setting and cannot be modified for individual clients.
interval	integer	If the number of warnings and failures over this interval value exceeds the configured limit for a client, the specified block action will be enforced. The default is 20. This is a global setting and cannot be modified for individual clients.
block unblock	<i>ip_address</i> [<i>minutes</i>]	Blocks a specific IP address for the number of minutes listed. If the optional <i>minutes</i> argument is omitted, the client is blocked until explicitly unblocked. Unblock releases a specific IP address.
clear-maximum		Clears all maximum statistics from the appliance.
default block-action	<i>drop</i> <i>send-tcp-rst</i>	Indicates the behavior when clients are at the maximum number of connections or exceed the warning limit: drop the connections that are over the limit or send TCP RST for connections over the limit. The default is drop. This limit can be modified on a per-client basis.
default connection-limit	integer	Indicates the number of simultaneous connections between 1 and 65535. The default is 100. This limit can be modified on a per-client basis.
default concurrent-request limit	integer	Indicates the maximum number of simultaneous requests that effective client IP sources (with <i>client.effective_address</i> policy) or explicit client IP sources (without <i>client.effective_address</i> policy) are allowed to make. The default value is unlimited. This limit can be applied on a per-client basis.

Table 70–1 Changing Global Defaults (Continued)

default failure-limit	integer	Indicates the maximum number of failed requests a client is allowed before the proxy starts issuing warnings. Default is 50. This limit can be modified on a per-client basis. By default, failed requests (with regard to attack detection) are defined as the following: • Connection failures (DNS lookup errors, connection refused, connection timed out, host unreachable, and so on) • 4xx (excluding 401 and 407) and 5xx HTTP response codes returned from the ProxySG or origin content server. • Each failure request event adds a count of one failure by default. The default definition for both the response code and the associated value per failed request event can be overridden via the CPL If the appliance serves an exception page to the client instead of serving a page returned by the server, the response code associated with the exception is used to decide if it was a failure or not.
default monitor-only		Enables monitor-only mode, which logs the defined thresholds that have been exceeded, but does not enforce the rules. The default value is disabled. This limit can be modified on a per-client basis. Note: The monitor-only mode setting has a higher precedence level than the default enforce mode. Enabling monitor-only mode disables rule enforcement.
no default monitor-only		Disables monitor-only mode. The default value is disabled. This limit can be modified on a per-client basis.
default request-limit	integer	Indicates the maximum number of HTTP requests that IP sources are allowed to make during a one-minute interval. The default value is unlimited. This limit can be applied on a per-client basis.
default unblock-time	minutes	Indicates the amount of time a client is locked out when the client-warning-limit is exceeded. By default, the client is blocked until explicitly unblocked. The default is unlimited. This limit can be modified on a per-client basis.
default warning-limit	integer	Indicates the number of warnings sent to the client before the client is blocked and the administrator is notified. The default is 10; the maximum is 100. This limit can be modified on a per-client basis.

To create and edit a client IP address:

Client attack-detection configuration is used to control the behavior of virus-infected machines behind the ProxySG.

1. Verify the system is in the attack-detection client submode.

```
SGOS#(config) attack-detection
SGOS#(config attack-detection) client
SGOS#(config client)
```

2. Create a client.

```
SGOS#(config client) create {ip_address | ip_prefix}
```

3. Move to edit client submode.

```
SGOS#(config client) edit client_ip_address
```

The prompt changes to:

```
SGOS#(config client ip_address)
```

4. Change the client limits as necessary.

```
SGOS#(config client ip_address) block-action drop | send-tcp-rst
SGOS#(config client ip_address) concurrent-request-limit
integer_between_1_and_2147483647
SGOS#(config client ip_address) connection-limit
integer_between_1_and_65534
SGOS#(config client ip_address) failure-limit
integer_between_1_and_500
SGOS#(config client ip_address) request-limit
integer_between_1_and_2147483647
SGOS#(config client ip_address) unblock-time
minutes_between_1_and_1440
SGOS#(config client ip_address) warning-limit
integer_between_1_and_100
```

Table 70–2 Changing the Client Limits

block-action	drop send-tcp-rst	Indicates the behavior when the client is at the maximum number of connections: drop the connections that are over the limit or send TCP RST for the connection over the limit. The default is drop.
concurrent-request-limit	<i>integer</i>	Indicates the maximum number of simultaneous requests that effective client IP sources (with <code>client.effective_address</code> policy) or explicit client IP sources (without <code>client.effective_address</code> policy) are allowed to make. The default value is unlimited.
connection-limit	<i>integer</i>	Indicates the number of simultaneous connections between 1 and 65534. The default is 100.
failure-limit	<i>integer</i>	Indicates the maximum number of failed requests a client is allowed before the proxy starts issuing warnings. The default is 50 and the maximum is 500.
monitor-only		Enables monitor-only mode, which logs the defined thresholds that have been exceeded, but does not enforce the rules. The default value is disabled. Note: The monitor-only mode setting has a higher precedence level than the default enforce mode. Enabling monitor-only mode disables rule enforcement.

Table 70-2 Changing the Client Limits (Continued)

request-limit	integer	Indicates the maximum number of HTTP requests that IP sources are allowed to make during a one-minute interval. The default value is unlimited. This limit can be applied on a per-client basis.
unblock-time	minutes	Indicates the amount of time a client is locked out when the client-warning-limit is exceeded. By default, the client is blocked until explicitly unblocked. The default is unlimited.
warning-limit	integer	Indicates the number of warnings sent to the client before the client is locked out and the administrator is notified. The default is 10; the maximum is 100.

To view the specified client configuration:

Enter the following command from the edit client submode:

```
SGOS#(config client ip_address) view
Client limits for 10.25.36.47:
Client concurrent request limit: unlimited
Client connection limit:          100
Client failure limit:             50
Client request limit:             unlimited
Client warning limit:             1
Blocked client action:            Drop
Client connection unblock time:  unlimited
Monitor only mode: disabled
```

To view the configuration for all clients:

1. Exit from the edit client submode:

```
SGOS#(config client ip_address) exit
```

2. Use the following syntax to view the client configuration:

```
view {<Enter> | blocked | connections | statistics}
```

To view all settings:

```
SGOS#(config client) view <Enter>
Client limits enabled:          true
Client interval:                20 minutes
Default client limits:
Client concurrent request limit: unlimited
Client connection limit:        100
Client failure limit:           50
Client request limit:           unlimited
Client warning limit:           1
Blocked client action:          Drop
Client connection unblock time: unlimited
Monitor only mode:              disabled
```

```
Client limits for 10.25.36.47:
  Client concurrent request limit: unlimited
  Client connection limit:        700
  Client failure limit:           50
  Client request limit:           unlimited
  Client warning limit:           1
  Blocked client action:          Drop
  Client connection unblock time: unlimited
  Monitor only mode:             disabled
```

To view the number of simultaneous connections to the ProxySG:

```
SGOS#(config client) view connections
Client IP      Connection Count
127.0.0.1      1
10.9.16.112    1
10.2.11.133    1
```

To view the number of blocked clients:

```
SGOS#(config client) view blocked
Client          Unblock time
10.11.12.13     2004-07-09 22:03:06+00:00UTC
10.9.44.73      Never
```

Note: The following thresholds dictate when a client receives a warning:

- Number of connections
- Number of failures
- Number of requests a client is allowed to make during a one-minute period
- Number of concurrent requests a client is allowed to make during a one-minute period

A client will receive a warning whenever a defined limit is exceeded by the client. If the client exceeds the configured warning limit, the client is then blocked.

To view client statistics:

```
SGOS#(config client) view statistics
Client IP  Failure Count  Warning Count  Request Count  Concurrent \
Request Count
10.9.44.72      1              1              0              0
```

To view specific maximum statistics for clients:

Enter the following syntax from the edit client submode to view the maximum statistics for a threshold category for a specified number of clients:

```
SGOS#(config client) view statistics maximum {requests | concurrent-
requests | failures | warnings | connections} <number_of_clients>
```

where: <number_of_clients> = an integer between 1 - 1024.

To disable attack-detection mode for all clients:

```
SGOS#(config client) disable-limits
```

To change the attack detection failure weight:

To change the default value of a single failed request event on the ProxySG appliance, you need to apply the **Set Attack Detection Failure** object. The object exists in the **Web Access Layer** as an **Action**. Each failed request can have a value of 0 - 500, depending on the nature of the failed request.

Note: Refer to the *Visual Policy Manager Reference* for complete details about the VPM.

To create attack detection failure weight policies:

1. Select the **Configuration > Policy > Visual Policy Manager** tab.
2. Click **Launch**. The VPM launches in a separate window.
3. Select **Policy > Add Web Access Layer**. An **Add New Layer** dialog displays.
4. Enter a name that is easily recognizable and click **OK**. A new policy tab and rule display in the VPM manager window.
5. Select **Action** under the new rule. Right click **Any > Set**. The **Set Action Object** window displays.
6. Select **New > Set Attack Detection** to add a new object.
7. The **Add Attack Detection Failure Object** window allows you to configure the attack detection weight value.
 - a. In the **Name** field, enter a name for the object or leave as is to accept the default.
 - b. From the **Failure Weight** field, enter an integer value between 0-500. This value is the amount by which the client's failure counter increases per failure event.
8. Click **OK**.
9. Click **OK** to return to the VPM.
10. Click the **Install Policy** button when finished adding policies.

To enforce ADP thresholds on the client's effective IP address:

If you rely on a deployment model where the client's real IP address is obscured by a load balancer or HTTP proxy, such as a reverse proxy indirect or forward proxy indirect deployment, you can configure ADP to use the value contained in the X-Forwarded-For header field or another custom header to identify the originating IP address. When clients have been identified using their effective client IP address, the specified thresholds which dictate when a client is blocked are applied.

To configure the appliance to extract the effective IP address from the request header, you need to specify the request header variable within policy. Keep in mind that the ProxySG appliance can only extract the effective IP address where so defined in the request header. If the request header is not present or is an invalid IP, the request will use the client IP instead.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file layers.

- ❑ Set the ProxySG appliance to extract the first IP address presented in the X-Forwarded-For header variable as the effective IP address.

```
<Proxy>
client.address=<ip_address> \
client.effective_address("${request.header.X-Forwarded-For}")
```

where:

<i>ip_address</i>	Specifies the HTTP proxy or load balancer IP address.
("\$(request.header.X-Forwarded-For)")	The effective IP address.

Notes

- Concurrent request limiting thresholds count requests from effective IP addresses (if `client.effective_address()` is present in policy) or explicit IP addresses (if `client.effective_address()` is not present policy) when using the `concurrent-request-limit` CLI command. `connection-limit` does not take effective IP clients into account and should not be used.
- Symantec recommends replacing all instances of `client.address` in existing policy with `client.effective_address` for all policies referencing the actual client IP instead of the IP of the downstream proxy or load balancer.

Creating the CPL

Be aware that the examples below are just part of a comprehensive authentication policy. By themselves, they are not adequate for your purposes.

Note: Refer to the *Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file layers.

- ❑ Set the failure weight value for a specific HTTP response code.

```
<proxy>
http.response.code=<CODE> attack_detection.failure_weight(<N>)
```

where:

<i>CODE</i>	HTTP Response Code	Specifies an HTTP response code to be defined as a failed request event.
N	<i>Failure weight</i>	Sets the failure weight value for the specified HTTP response code per failed request event. If set to 0, the response code is not counted as a failure.

Configuring Attack-Detection Mode for a Server or Server Group

Server attack-detection configuration is used when an administrator wants to protect a server from becoming overloaded by too many active connections.

You can create, edit, or delete a server. A server must be created before it can be edited. You can treat the server as an individual host or you can add other servers, creating a server group. All servers in the group have the same attack-detection parameters, meaning that if any server in the group gets the maximum number of simultaneous requests, all new connections to the servers in the group are blocked.

You must create a server group before you can make changes to the configuration.

To create a server or server group:

1. At the (config) prompt:

```
SGOS#(config) attack-detection
SGOS#(config attack-detection) server
```

The prompt changes to:

```
SGOS#(config server)
```

2. Create the first host in a server group, using the fully qualified domain name:

```
SGOS#(config server) create hostname
```

To edit a server or server group:

At the (config server) prompt:

```
SGOS#(config server) edit hostname
```

The prompt changes to (config server hostname).

```
SGOS#(config server hostname) {add | remove} hostname
SGOS#(config server hostname) concurrent-request-limit
integer_from_1_to_65535
```

where:

<i>hostname</i>		The name of a previously created server or server group. When adding a hostname to the group, the hostname does not have to be created. The host that was added when creating the group cannot be removed.
add remove	<i>hostname</i>	Adds or removes a server from this server group.
concurrent-request-limit	<i>integer</i>	Indicates the number of simultaneous requests allowed from this server or server group. The default is 1000.

To view the server or server group configuration:

```
SGOS#(config server hostname) view
Server limits for hostname:
concurrent-request limit: 1500
```


Chapter 71: Authenticating a ProxySG

This section describes device authentication, which is a mechanism that allows devices to verify each others' identity; devices that are authenticated can be configured to trust only other authenticated devices.

Device authentication is important in several situations:

- ❑ Securing the network. Devices that are authenticated have exchanged certification information, verified each others' identity and know which devices are trusted.
- ❑ Securing protocols. Many protocols require authentication at each end of the connection before they are considered secure.

Note: ProxySG authentication is always used in association with other SGOS features. For example, you can use appliance authentication with the ADN implementation of secure tunnels. The secure tunnels feature uses authentication, the process of verifying a device's identity, with authorization, the process of verifying the permissions that a device has. For information on secure tunnels and appliance authentication, see ["Securing the ADN"](#) on page 739.

This section includes the following topics:

- ❑ ["ProxySG Device Authentication Overview"](#)
- ❑ ["Appliance Certificates and SSL Device Profiles"](#) on page 1290
- ❑ ["Obtaining a ProxySG Appliance Certificate"](#) on page 1292
- ❑ ["Obtaining a Non-Blue Coat Appliance Certificate"](#) on page 1295
- ❑ ["Creating an SSL Device Profile for Device Authentication"](#) on page 1296

ProxySG Device Authentication Overview

The Blue Coat implementation allows devices to be authenticated without sending passwords over the network. Instead, a device is authenticated through certificates and SSL device profiles that reference the certificates. Both the profile and the referenced certificate are required for device authentication.

- ❑ **Certificates:** Certificates contain information about a specific device. Blue Coat runs an Internet-accessible Certificate Authority (CA) for the purpose of issuing appliance certificates to SGOS devices. You can also create your own appliance certificates.
- ❑ **Profiles:** A profile is a collection of information used for several purposes, such as device-to-device authentication or when the ProxySG is an SSL endpoint for non-proxy traffic.

The ProxySG comes with three built-in profiles: *bluecoat-appliance-certificate*, *default*, and *passive-attack-protection-only*. A profile can indicate whether the device has a certificate and if the certificates of other devices should be verified. You can create other profiles to change the default settings. The *bluecoat-appliance-certificate* profile is the one that is used for device authentication; this profile references the appliance certificate on your ProxySG.

Appliance Certificates and SSL Device Profiles

In the Blue Coat implementation of device authentication, both an appliance certificate and an SSL device profile that references the appliance certificate keyring are required for device authentication to be successful. Each device to be authenticated must have an appliance certificate and a profile that references that certificate.

Note that device authentication does not take effect unless the SSL device profile is enabled; for example, if you use WAN optimization, you enable the profile on the **Configuration > ADN > General > Device Security** tab.

About ProxySG Appliance Certificates

ProxySG hardware appliances come with a cryptographic key that allows the system to be authenticated as an ProxySG appliance when an *appliance certificate* is obtained. Note that appliance certificates are not relevant in a virtual machine environment.

An appliance certificate is an X.509 certificate that contains the hardware serial number of a specific ProxySG as the CommonName (CN) in the subject field. This certificate then can be used to authenticate the ProxySG appliance whose hardware serial number is listed in the certificate. Information from the presented certificate is extracted and used as the *device ID*.

Blue Coat runs an Internet-accessible CA for the purpose of issuing appliance certificates. The root certificate for the Blue Coat CA is automatically trusted by SGOS for device authentication. These Blue Coat-signed certificates contain no authorization information and are valid for five years.

You can provide your own device authentication certificates for the ProxySG appliances on your network if you prefer not to use the Blue Coat CA.

About SSL Device Profiles

An SSL device profile contains the information required for device authentication:

- ❑ The name of the keyring that contains the private key and certificate this device uses to authenticate itself. The default keyring is `appliance-key`. (For information on private and public keys, see ["Public Keys and Private Keys"](#) on page 1114.)
- ❑ The name of the CA Certificate List (CCL) that contains the names of certificates of CAs trusted by this profile. If another device offers a valid certificate signed by an authority in this list, the certificate is accepted. The default is `appliance-ccl`. For information on CCLs, see ["Managing CA Certificate Lists"](#) on page 1143.
- ❑ Verification of the peer certificate.

When the ProxySG is participating in device authentication as an SSL client, the peer certificate verification option controls whether the server certificate is validated against the CCL. If verification is disabled, the CCL is ignored.

When the ProxySG is participating in device authentication as an SSL server, the peer certificate verification option controls whether to require a client certificate. If verification is disabled, no client certificate is obtained during the SSL handshake. The default is `verify-peer-certificate enabled`.

- ❑ Specification of how the device ID authorization data is extracted from the certificate. The default is `$(subject.CN)`.
- ❑ SSL cipher settings. The default is AES256-SHA.

Each Blue Coat appliance has an automatically-constructed profile called **bluecoat-appliance-certificate** that can be used for device-to-device authentication. This profile cannot be deleted or edited.

If you cannot use the built-in profile because, for example, you require a different cipher suite or you are using your own appliance certificates, you must create a different profile, and have that profile reference the keyring that contains your certificate.

Note: If you do not want to use peer verification, you can use the built-in **passive-attack-detection-only** profile in place of the **bluecoat-appliance-certificate** profile.

This profile uses a self-signed certificate and disables the `verify-peer` option, so that no authentication is done on the endpoints of the connection. The traffic is encrypted, but is vulnerable to active attacks.

This profile can be used only when there is no threat of an active man-in-the-middle attack. Like the **bluecoat-appliance certificate** profile, the **passive-attack-detection-only** profile cannot be edited or deleted.

If you create your own profile, it must contain the same kind of information that is contained in the Blue Coat profile. To create your own profile, skip to ["Creating an SSL Device Profile for Device Authentication"](#) on page 1296.

Obtaining a ProxySG Appliance Certificate

In many cases, if you have Internet connectivity, an appliance certificate is automatically fetched by the ProxySG, and no human intervention is required. In other cases, if the Internet connection is delayed or if you do not have Internet access, you might have to manually initiate the process of obtaining an appliance certificate.

How you obtain an appliance certificate depends upon your environment:

- ❑ If the device to be authenticated has Internet connectivity and can reach the Blue Coat CA server, continue with "[Automatically Obtaining an Appliance Certificate](#)" on page 1292.
- ❑ If the device to be authenticated cannot reach the Blue Coat CA server, you must acquire the certificate manually; continue with "[Manually Obtaining an Appliance Certificate](#)" on page 1292.

Important: Appliance certificates are not relevant in a virtual machine environment.

Automatically Obtaining an Appliance Certificate

The appliance attempts to get the certificate completely automatically (with no user intervention) if it can connect to the Blue Coat CA server at boot time or within about five minutes of being booted. If the appliance does not have a certificate (for example, it had one until you did a `restore-defaults factory-defaults` command) it attempts to get one on every boot. Once the appliance gets a certificate, that certificate is used until another `restore-defaults factory-defaults` command is issued.

If Internet connectivity is established more than five minutes after the system is booted, you might need to complete the following steps.

To automatically obtain an appliance certificate:

1. Select the **Configuration > SSL > Appliance Certificates > Request Certificate** tab.
2. Click **Request appliance certificate**.

The Blue Coat CA server does validation checks and signs the certificate. The certificate is automatically placed in the `appliance-key` keyring. Note that the `appliance-key` keyring cannot be backed up. The keyring is re-created if it is missing at boot time.

Manually Obtaining an Appliance Certificate

Complete the following steps to obtain an appliance certificate manually. The overview of the procedure is to:

- ❑ Generate a appliance certificate signing request and send it to the Blue Coat CA server for verification and signature.
- ❑ Import the signed certificate into the ProxySG.

To generate a CSR:

1. Select the **Configuration > SSL > Appliance Certificates > Request Certificate** tab.
2. Select **Create CSR**. The Appliance Certificate Signing Request dialog displays.



3. Copy the certificate request, including the certificate request signature. Be sure to include the **Begin Certificate** and **End Certificate** statements, as well as the **Begin CSR Signature** and **End CSR Signature** statements.
4. Click **OK**.
5. Go to the Blue Coat CA server Web site at <https://abrca.bluecoat.com/sign-manual/index.html>.

Blue Coat - ABRCA Manual Form	
CSR+Signature	<pre> -----BEGIN CERTIFICATE REQUEST----- MIIBxzCCATACAQAwgYYxCzAJBgNVBAYTA1VTMQswCQYDVQQIEwJQTESMBAGA1UE BxMJU3Vubn12YWx1MSAwHgYDVQQKEXdCbHV1IENvYXQGU31zdGVtcywgSW5jLjEf MBOGA1UECzMWQmx1ZSBDb2FOIFNHNHJAwIFN1cm1lc3ETMBEGA1UEAxMKNDUwNTA2 MDAyMDCBnzANBgkqhkiG9w0BAQEFAAOBjQAwgYkCgYEAAt1Y3SaIts6ADY8BPYhVb MrvA/aRsOEd6OCwswuxiFSTHn3ijDug5ttT5DDvfmxy4YcgP7vdTaeVdqeQDBwkM FpjxxWfjTHXIngeBWKThM8WlcO/gBa8z7cWUryxyGS9FS3H2ZBZXsXsVQT19zWu3 2XA3QtIir8RH7MM1dbPomrOCAwEAaAaAAMAGCSqGS1b3DQEBBAUAA4GBAHBikVoc TjM8zDmPttII7dMNCwmfPIy3zeyppdrMFL1JcnJjwqh1XrndN7WHYUXEwhYJtU9p 70kyFs+giBtIzdd8fnZaeF4JXNCzSfLqWnpKOjTBA9WLTMc1Th1Hp1UZE/T11DRS kZ6AfyhJQGhxmKuAi8LLRjPMO5Y0owxo8A17 -----END CERTIFICATE REQUEST----- -----BEGIN CSR SIGNATURE ----- KsrqFGa5jb2Az+GL/Hm90FmmBzLg0svAwBbaYD64qNm3VH17ADaMw2LfrZ1D13ez BOgXKJEBU5w7TULG23QJV3XpwP7XOb6ms1ekg/XPNZ2OmoNjI3VreJ+A9usYpUhh 56qFKfIcivndchukrhI= -----END CSR SIGNATURE ----- </pre>

- Paste the CSR and signature into the CSR panel.
- Click **Generate Cert.**

The signed certificate displays, and can be pasted into the appliance-key keyring.

```

-----BEGIN CERTIFICATE-----
MIIF/jCCBOagAwIBAgICAMowDQYJKoZIhvcNAQEFBQAwwgYxChAJBgNVBAYTA1VT
MRMwEQYDVQQIEwZlYXN1Zm9udm91dC5jYXN1Zm9udm91dC5jYXN1Zm9udm91dC5j
BAoTF0JsdWUgQ29hdCBTeXN0ZW1zLCBjb250bWw0bWw0bWw0bWw0bWw0bWw0bWw0
IEFCUkNBMRswGQYDVQQDEwJhbn12YWx1MSAwHgYDVQQKEwJQTESMBAGA1UECzMW
CQEWFXN5c2FkbWluQGJsdWVjb2F0LmNvbTAeFw0wNzAxMjkyMDM5NDdaFw0xMjAx
MjkyMDM5NDdaMIGGMQswCQYDVQQGEwJVUzELMAkGA1UECBMCQ0ExEjAQBGNVBACt
CVN1bm55dmFsZTEgMB4GA1UEChMXQmx1ZSBDb2FOIFN5c3RlbXMsIEluYy4xHzAd
BgNVBAsTFkJsZWUgQ29hdCBTRzIwMCM5MjkyMDM5NDdaFw0xMjAxMjkyMDM5NDda
OTIwZj8wDQYJKoZIhvcNAQEBBQADgY0AMIGJAoGBAMBUmCuKSSsd+D5kJQiWu3OG
DNLCvf7SyKK5+SBCJU2iKwP5+EfiQ5JsScWJghtIo94EhdSC2zvBPQqWbZAJXN74
k/yM4w9ufjfo+G7xPYcMrGmwVBGnXbEhQkagclFH2orINNY8SVDYVL1V4dRM+0at
YpEiBmSxipmRSMZL4kqtAgMBAAGjggLGMICwJAJBgNVHRMEAjAAMAsGA1UdDwQE
AwIE8DBOBgNVHSEUerzBFBggrBgEFBQcDAQYIKwYBBQUHAWIGCCsGAQUFBwMEBgsr
BgEEAfElAQECAQYLKwYBBAHxJQEBAgIGCysGAQQB8SUBAQIDMB0GA1UdDgQWBBSF
NqC2ubTI7OT5j+KqCPGLSDO7DzCB6wYDVROjBIHjMIHggBSwEYwcq1N6G1ZhpCxn
OTIu8fNe1aGBvKSBUtCBtjELMAkGA1UEBhMCVVMxZARBgNVBAGTCkNhbg1mb3Ju
aWEwEjAQBGNVBACtCVN1bm55dmFsZTEgMB4GA1UEChMXQmx1ZSBDb2FOIFN5c3Rl
bXMsIEluYy4xGTAXBGNVBAsTEEJsdWUgQ29hdCwgQUJSQ0ExGzAZBgNVBAMTEmFi

```

```
cmNhLmJsdWVjb2F0LmNvbTEkMCIGCSqGSIB3DQEJARYVc3lzYWRTaW5AYmx1ZWNV
YXQuY29tggkAhmhBUPEEb60wgZ8GCCsGAQUFBwEBBIGSMIGPMekGCCsGAQUFBzAB
hjlodHRwczoL2FicmNhLmJsdWVjb2F0LmNvbS9jZ2ktYmluL2Rldm1jZS1hdXR0
ZW50aWNhdGlvbi9vY3NwMEIGCCsGAQUFBzAChjZodHRwOi8vYWJyY2EuYmx1ZWNV
YXQuY29tL2Rldm1jZS1hdXR0ZW50aWNhdGlvbi9jYS5jZ2kwSAYDVR0fBEEwPzA9
oDugOYY3aHR0cDovL2FicmNhLmJsdWVjb2F0LmNvbS9kZXZpY2UtYXV0aGVudGlj
YXRpb24vQ1JMLmNybDBfBgNVHSAEWDBMWfQGCisGAQQB8SUBAQEwRjBEBggrBgEF
BQcCARY4aHR0cDovL2FicmNhLmJsdWVjb2F0LmNvbS9kZXZpY2UtYXV0aGVudGlj
YXRpb24vcnBhLmh0bWwwDQYJKoZIhvcNAQEFBQADggEBACIhQ7Vu6aGJBpxP255X
d2/Qw7NiVsnqOlAy913QZlieFfVATJnCeSrH+M9B/2XtnRxVT0/ZWrf4GbsdYqTF
hc9jR/IwKu6kZq32Dqo8qFU5OzbAEzT2oebB5QgwuJtHcJHggp9PS9uS27qAnGQK
OeB2bYcjWtMvTvr50iDOV69BEQz+VXos8QiZmRHLVnebQSjl3bilw3VjBw31tCmc
clgz0SlN9ZmJdRU/PlWdNVqD40LqcMZQ53HqcdWNEzN2uvigIb//rM7XazK7xIaq
r23/+BsZlYKAeVMq3PEmxA2zLzO+jf79a8ZvIKrF27nNuTN7NhFL/V6pWNE1o9A
rbs=
-----END CERTIFICATE-----
```

To import a certificate onto the ProxySG appliance:

1. Copy the certificate to your clipboard. Be sure to include the `Begin Certificate` and `End Certificate` statements.
2. Select the **Configuration > SSL > Keyrings** tab.
3. Select the keyring that is used for device authentication. The keyring used by the **bluecoat-appliance-certificate** profile is the `appliance-key` keyring.
4. Click **Edit** in the **Keyrings** tab.
5. In the **Certificate** panel, click **Import**.
6. Paste the certificate you copied into the dialog box.
7. Click **Close**.

Obtaining a Non-Blue Coat Appliance Certificate

If you use your own CA to create certificates for device authentication, complete the following steps:

1. Create a keyring for the appliance's certificate. For information on creating a keyring, see ["Creating a Keyring"](#) on page 1119.
2. Generate the certificate signing request and get it signed. For information on creating a CSR, see ["Creating a CSR"](#) on page 1130.

Note: You cannot put a Blue Coat appliance certificate into a keyring you create yourself.

3. Create a CA certificate list. For information on creating a CCL, see ["Managing CA Certificate Lists"](#) on page 1143.
 - a. Import the CA's root certificate.
 - b. Add the certificate to the CCL.
4. Create a device profile. For information on creating a profile, see ["Appliance Certificates and SSL Device Profiles"](#) on page 1290.
5. Associate the device profile with the keyring and CCL. The keyring and CCL must already exist.
6. Adjust other parameters, including authorization data extractor (if the certificate is to be used for authorization), as needed.
7. Configure each application that uses device authentication to reference the newly created profile.

For more information, see ["About SSL Device Profiles"](#) on page 1170.

Creating an SSL Device Profile for Device Authentication

An SSL device profile only needs to be created if you cannot use the built-in **bluecoat-appliance-certificate** profile without modification; note that the **bluecoat-appliance-certificate** profile cannot be deleted or edited.

Additional profiles with different settings can be created; for example, if you require a different cipher setting than what the **bluecoat-appliance-certificate** profile uses, you can create a profile with the different cipher suite.

To create or edit an SSL device profile:

1. Select the **Configuration > SSL > Device Profiles > Profiles** tab.
2. Click **New**.

Create SSL Device Profile

Create SSL Device Profile

Name:

Usable for: Unknown

SSL protocols: ☒ TLSv1.2 ☒ TLSv1.1 ☒ TLSv1 ☐ SSLv3* ☐ SSLv2*

* These SSL protocols are not recommended.

Keyring:

CCL:

Device ID extractor: ☒ Verify peer

Device ID:


Selected Ciphers: AES128-SHA256 AES256-SHA256 AES128-SHA AES256-SHA DHE-RSA-AES128-SHA DHE-RSA-AES256-SHA DES-CBC3-SHA RC4-SHA RC4-MD5 DES-CBC-SHA EXP-DES-CBC-SHA EXP-RC4-MD5 EXP-RC2-CBC-MD5

3. **Name:** Give the profile a meaningful name. (If you are editing the default profile, this field is grayed out.) The only valid characters are alphanumeric, the underscore, and hyphen, and the first character must be a letter.
4. **SSL protocol versions:** Change the default from **TLS1.2**, **TLS1.1**, **TLSv1** to any other protocol listed as required.
5. **Keyring:** If the server in question requires a client certificate, then select the keyring used to negotiate with origin content servers through an encrypted connection. Only keyrings with certificates can be associated with the SSL client, displayed in the **Keyring** drop-down list. By default, no keyring is selected.

Note: You must create a new keyring for device authentication if you do not use the `appliance-key` keyring. The other keyrings shipped with the ProxySG are dedicated to other purposes. For information on creating a new keyring, see ["Creating a Keyring"](#) on page 1119.

6. **CCL:** From the drop-down list, select the CA Certificate List you want to use. The **browser-trusted** CCL is the default.
7. **Device ID extractor:** The field describes how device ID information is extracted from a presented certificate. The string contains references to the attributes of the subject or issuer in the form `$(subject.attr[.n])` or `$(issuer.attr[.n])`, where `attr` is the short-form name of the attribute and `n` is the ordinal instance of that attribute, counting from 1 when the subject is in LDAP (RFC 2253) order. If `n` is omitted, it is assumed to be 1.

The default is `$(subject.CN)`; many other subject attributes are recognized, among them `OU`, `O`, `L`, `ST`, `C`, and `DC`.
8. **Verify peer:** This setting determines whether a client certificate is requested on incoming SSL connections and verified against the specified CCL. This is enabled by default.



9. **Selected ciphers:** To use a different cipher suite:
 - a. Click **Edit Ciphers**. A dialog displays a list of cipher suites. The cipher suites available for use depend on the protocols you selected. For improved security, select only ciphers with HIGH strength.
 - b. To add ciphers to the list, select them from the list of available cipher suites and click **Add**. To remove selected ciphers, select them and click **Remove**.

Note: Because they contain known vulnerabilities, Symantec recommends that you do not use the SSLv3 and SSLv2 protocols; however, if you do select the SSLv2 protocol, additional cipher suites are available: DES-CBC3-MD5 (High, 168-bit), RC2-CBC-MD5 (Medium, 128-bit), and DES-CBC-MD5 (Low, 56-bit).

10. Click **OK** to close the dialog.
11. Click **Apply**.

Chapter 72: Monitoring the ProxySG

This section describes the methods you can use to monitor your ProxySG appliances, including disk management, event logging, monitoring network devices (SNMP), and health monitoring. The section also provides a brief introduction to Director.

Topics

- ❑ [Section A: "Using Director to Manage ProxySG Systems" on page 1300](#)
- ❑ [Section B: "Monitoring the System and Disks" on page 1304](#)
- ❑ [Section C: "Configuring Event Logging and Notification" on page 1309](#)
- ❑ [Section D: "Monitoring Network Devices \(SNMP\)" on page 1319](#)
- ❑ [Section E: "Configuring Health Monitoring" on page 1335](#)

Section A: Using Director to Manage ProxySG Systems

Blue Coat Director allows you to manage multiple ProxySG appliances, simplifying configuration and setup and giving you a central management solution. You can configure one ProxySG appliance and use it as a template to configure other devices in a similar, or identical, way.

Other advantages of using Director include:

- ❑ Reducing management costs by centrally managing all ProxySG appliances.
- ❑ Eliminating the need to manually configure each ProxySG.
- ❑ Recovering from system problems with configuration snapshots and recovery.
- ❑ Monitoring the health of individual appliances or groups of appliances.

This section discusses the following topics:

- ❑ ["Automatically Registering the ProxySG with Director"](#)
- ❑ ["Setting Up SSH-RSA Without Registration"](#) on page 1303

Automatically Registering the ProxySG with Director

Director manages ProxySG appliances after you perform any of the following:

- ❑ *Register* the appliances with Director.

Registering an appliance with Director creates a secure connection using RSA-SSH (public/private key cryptography). During the registration process, Director replaces the following with values known only to Director:

- Appliance's administrative password
- Appliance's enable mode password
- Appliance's serial console password
- Front panel PIN

This is useful if you want to control access to the appliance or if you want to ensure that appliances receive the same configuration.

During registration, the ProxySG uses its Blue Coat appliance certificate or a registration password configured on Director to confirm identities before exchanging public keys. If the ProxySG has an appliance certificate, that certificate is used to authenticate the ProxySG to Director as an SSL client.

If the appliance does not have an appliance certificate, you must configure a registration password on Director and specify that password when you register the ProxySG. Refer to the *Symantec Director Configuration and Management Guide* for more information about specifying the shared secret.

- ❑ Manually *add* the appliances to Director.

Initially, SSH-Simple (user name/password) is used to authenticate the device with Director. You have the option of changing the authentication mechanism to SSH-RSA at a later time.

Note:

- Regardless of whether or not you register the appliance with Director, communication between the ProxySG appliance and Director is secured using SSHv2.
- The ProxySG uses interface 0:0 to register with Director. Before you attempt to register a ProxySG with Director, make sure its interfaces, static routes, and Internet gateways are configured properly to allow communication to succeed.
- The Blue Coat appliance certificate is an X.509 certificate that contains the hardware serial number of a specific ProxySG as the Common Name (CN) in the subject field. See ["Appliance Certificates and SSL Device Profiles"](#) on page 1290 for more information about appliance certificates.

Continue with one of the following sections:

- ["Registration Requirements"](#)
- ["Registering the ProxySG with Director"](#) on page 1301

Registration Requirements

To register the appliance with Director, the SSH Console management service on the ProxySG must be enabled. Director registration will fail if the SSH Console has been disabled or deleted, or if the SSHv2 host key has been deleted.

Ports 8085 and 8086 are used for registration from the ProxySG to Director. If Director is already in the network, you do not need to open these ports. If you have a firewall between the ProxySG and Director and you want to use the registration feature, you must open ports 8085 and 8086.

Continue with ["Registering the ProxySG with Director"](#).

Registering the ProxySG with Director

Though usually initiated at startup (with the serial console setup), you can also configure Director registration from the ProxySG's Management Console, as described in the following procedure.

For more information about registration, see one of the following sections:

- ["Automatically Registering the ProxySG with Director"](#) on page 1300
- ["Registration Requirements"](#) on page 1301

To register the appliance with a Director:

1. Select the **Maintenance > Director Registration** tab.

The screenshot shows the 'Register' tab with the following fields and buttons:

- Director IP address:** 192.0.2.45
- Director serial number:** 123456789
- Appliance name:** CorpTower25G3
- Buttons:** 'Retrieve S/N from Director' and 'Register'.

Numbered arrows indicate the registration steps: 2 points to the IP address field, 3 points to the serial number field, 4 points to the appliance name field, and 6 points to the Register button.

2. In the **Director IP address** field, enter the Director IP address.
3. In the **Director serial number** field, enter the Director serial number or click **RetrieveS/N from Director** (which is also a quick to verify that you entered a valid IP address in Step 2). If you retrieve the serial number from the Director, verify that the serial number matches the one specified for your Director.
4. (Optional) In the **Appliance name** field, enter a *friendly* name to identify the ProxySG.
5. If your appliance does not have an appliance certificate, enter the registration password in the **Registration password** field. (This field displays only if the appliance has no certificate.)

Note: Refer to the *Symantec Director Configuration and Management Guide* for more information about configuring the registration password. For information about appliance certificates, see [Chapter 61: "Managing X.509 Certificates"](#) on page 1113.

6. Click **Register**.

After the registration process is complete, Director communicates with the ProxySG using SSH-RSA. The appliance's administrative password, enable mode password, serial console password, and front panel PIN are values known only to Director.

Note: To verify or confirm that a ProxySG is registered with a Director (in the CLI):

```
#sh ssh-console director-client-key
```

This returns either:

```
No Director client key list installed
```

or

```
director xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx
```

For more information, see the *Blue Coat Director Configuration and Management Guide*.

Setting Up SSH-RSA Without Registration

If you manually add a device to Director, the authentication mechanism is SSH-Simple, meaning the appliance's user name and password are sent over the network as plain text. To securely authenticate the device with Director using SSH-RSA, you must do either of the following:

- ❑ (Recommended.) Push SSH-RSA keys to the device using the Director Management Console or command line. For more information, see the *Blue Coat Director Configuration and Management Guide*.
- ❑ Use the `import-director-client-key` CLI command from the ProxySG.

Complete the following steps to put Director's public key on the ProxySG using the CLI of the appliance. You must complete this procedure from the CLI. The Management Console is not available.

- a. Log in to the ProxySG you want to manage from Director.
- b. From the `(config)` prompt, enter the `ssh-console` submode:

```
SGOS#(config) ssh-console
SGOS#(config ssh-console)
```

- c. Import Director's key that was previously created on Director and copied to the clipboard.

Important: You must add the Director identification at the end of the client key. The example shows the username, IP address, and MAC address of Director. **Director** must be the username, allowing you access to passwords in clear text.

```
SGOS#(config services ssh-console) inline director-client-key
Paste client key here, end with "..." (three periods)
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEAvJIXt1ZausE9qrcXem2IK/
mC4dY8Cxxo1/B8th4KvedFY33OByO/pvwucuhPZz+b1LETTY/
zc3SL7jdVffq00KBN/ir4zu7L2XT68ML20RWa9tXFedNmKl/iagI3/
QQJ8T8zQM6o7WnBzTvMC/ZElMZZddAE3yPCv9+s2TR/
Ipk=director@10.25.36.47-2.00e0.8105.d46b
...
ok
```

To view the fingerprint of the key:

```
SGOS#(config ssh-console) view director-client-key clientID
jsmith@granite.example.com
83:C0:0D:57:CC:24:36:09:C3:42:B7:86:35:AC:D6:47
```

To delete a key:

```
SGOS#(config ssh-console) delete director-client-key clientID
```

Section B: Monitoring the System and Disks

The **System and disks** page in the Management Console has the following tabs:

❑ **Summary**

Provides configuration information and a general status information about the device.

❑ **Tasks**

Enables you to perform systems tasks, such as restarting the system and clearing the DNS or object cache. See "[Performing Maintenance Tasks](#)" on page 1399 for information about these tasks.

❑ **Environment**

Displays hardware statistics.

❑ **Disks**

Displays details about the installed disks and enables you take them offline.

❑ **SSL Cards**

Displays details about any installed SSL cards.

These statistics are also available in the CLI.

Note: The Management Console for ProxySG400 appliances does not contain an **Environment** tab.

System Configuration Summary

To view the system configuration summary, select **Maintenance > System and Disks > Summary**.

Summary	Tasks	Environment
Configuration		
Model:	210	
Disks installed:	1	
Memory installed:	1024 megabytes	
CPUs installed:	1	
IP address:	10.9.59.65	
Software version:	SGOS 5.4.1.1 MACH5 Edition	
Software release ID:	39730	
NIC 0 MAC:	00D083050780	
Serial number:	0408063430	
General Status		
System started:	2009-06-01 10:42:40-07:00PDT	
CPU utilization:	1 percent	

❑ **Configuration area:**

- **Model**—The model number of this ProxySG.
- **Disks Installed**—The number of disk drives installed in the ProxySG. The Disks tab displays the status of each drive.
- **Memory installed**—The amount of RAM installed in the ProxySG.
- **CPUs installed**—The number of CPUs installed in the ProxySG.
- **IP Address**—The IP address assigned to this ProxySG.
- **Software version**—The SGOS image name and edition type (Mach 5 or Proxy).
- **Serial release ID**—The SGOS image version number.
- **NIC 0 MAC**—The MAC address assigned to the connected interface(s).
- **Serial number**—The ProxySG serial number.

❑ **General Status area:**

- **System started**—The most recent time and date that the ProxySG was started.
- **CPU utilization**—The current percent of CPU usage.

Viewing System Environment Sensors

The icons on the Environment tab are green when the related hardware environment is within acceptable parameters and red when an out-of-tolerance condition exists. If an icon is red, click **View Sensors** to view detailed sensor statistics to learn more about the out-of-tolerance condition.


Note: The health monitoring metrics on the **Statistics > Health** page also display the state of environmental sensors. See [Section E: "Configuring Health Monitoring"](#) on page 1335 for more information.

The ProxySG400 model ProxySG does not support viewing environmental statistics.

To view the system environment statistics:

1. Select the **Maintenance > System and disks > Environment** tab (there might be a slight delay displaying this page as the system gathers the information).

Note: This displayed contents of this tab varies depending on the type of ProxySG. Systems with multiple disks display environmental information for each disk.



2. Click **View Sensors** to see detailed sensor values.

Sensor statistics		
Sensor Name	Reading	Status
MB Temperature	31.0 C	OK
CPU Temperature	31.0 C	OK

If any disk statistics display statuses other than **OK**, the ProxySG is experiencing environmental stress, such as higher than advised heat. Ensure the area is properly ventilated.

Viewing Disk Status and Taking Disks Offline


You can view the status of each of the disks in the system and take a disk offline if needed.

To view disk status or take a disk offline:

1. Select the **Maintenance > System and disks > Disks 1-2** tab.


The default view provides information about the disk in slot 1.

Note: The name and displayed contents of this tab differs, depending on the range of disks available to the ProxySG model you use.



Information displays for each present disk.

2. (Optional) To take a disk offline:
 - a. Select a disk and click the **Take disk X offline** button (where *x* is the number of the disk you have selected). The Take Disk Offline dialog displays.



- b. Click **OK**.

Note: Since there are no physical appliance disks in a virtual appliance, the **Take disk X offline** button is not available on the ProxySG VA.

Viewing SSL Accelerator Card Information

Selecting the **Maintenance > System and disks > SSL Cards** tab allows you to view information about any SSL accelerator cards in the system. If no accelerator cards are installed, that information is stated on the pane.

Note: You cannot view statistics about SSL accelerator cards through the CLI.

To view SSL accelerator cards:

Select the **Maintenance > System and disks > SSL Cards** tab.

Summary	Tasks	Environment	Disks 1-2	SSL Cards
- SSL Accelerators				
Internal: Cavium CN501 Security Processor				

Section C: Configuring Event Logging and Notification

You can configure the ProxySG appliance to log system events as they occur. The event logging options enable you to:

- ❑ Select the event logging levels to view in the event log.
See ["Selecting Which Events to View"](#) on page 1309.
- ❑ Configure the size of the event log.
See ["Setting Event Log Size"](#) on page 1310.
- ❑ Configure notification for events.
See ["Enabling Event Notification"](#) on page 1310.
- ❑ Configure Syslog monitoring.
See ["Syslog Event Monitoring"](#) on page 1312.
- ❑ View the event log.
["Viewing Event Log Configuration and Content"](#) on page 1316


Selecting Which Events to View

All events are logged to the event log. However, the events displayed in the event log (when viewed from the CLI, MC, or syslog) correspond to the configured event logging levels. Select an event logging level to view; deselect to omit. When you select an event level, all levels above the selection are included. For example, if you select **Verbose**, all event levels are included.

The event logging level options are listed from the most to least important events.

To set the event logging level:

1. Select the **Maintenance > Event Logging > Level** tab.



2. Select the events to log:

<i>Severe errors</i>	Displays only severe error messages in results.
<i>Configuration events</i>	Displays severe and configuration change error messages in results.

<i>Policy messages</i>	Displays severe, configuration change, and policy event error messages in results.
<i>Informational</i>	Displays severe, configuration change, policy event, and information error messages in results.
<i>Verbose</i>	Displays all error messages in results.

3. Click **Apply**.

Setting Event Log Size

You can limit the size of the appliances's event log and specify what occurs when the log size limit is reached.

To set event log size:

1. Select the **Maintenance > Event Logging > Size** tab.

The screenshot shows the 'Size' tab in the event logging configuration. It includes a section for 'Event log size' with a text input for 'Limit event log to' set to '10' megabytes of disk space. Below this is a section 'When event log reaches maximum size:' with two radio button options: 'Overwrite earlier events' (which is selected) and 'Stop logging new events'.

2. In the **Event log size** field, enter the maximum size of the event log in megabytes. The default is 10 MB.
3. Select the action that occurs when the event log reaches maximum size:
 - **Overwrite earlier events**—The ProxySG overwrites the older half of the event entries, replacing it with the most recent events. There is no way to recover the overwritten events.
 - **Stop logging new events**—The ProxySG retains all of the entries to date, but new events are not recorded.
4. Click **Apply**.

Enabling Event Notification

You can configure the ProxySG appliance to send event notifications to individuals in your organization using SMTP. To do this, you require either the hostname or IP address and port of the SMTP server, as well as the email addresses of the recipients.

You can also send event notifications directly to Symantec for support purposes. The Symantec SMTP gateway sends mail only to Symantec; it does not forward mail to other domains. For information on configuring diagnostic reporting, see [Chapter 75: "Diagnostics"](#).

To enable event notifications:

1. Select **Maintenance > Event Logging > Mail**.

The screenshot shows the 'Mail' configuration page. At the top are tabs for 'Level', 'Size', 'Mail', and 'Syslog'. The 'Mail' tab is selected. Below the tabs is a section titled 'Mail notifications to:' which contains a list of email addresses. The first address, 'hub.porter@company.com', is highlighted. To the right of the list are three buttons: 'New', 'Edit', and 'Delete'. Below these buttons are two radio buttons: 'SMTP gateway name' (which is selected) and 'SMTP gateway IP'. To the right of the 'SMTP gateway name' radio button is a text field containing 'mail.server.company.com'. Below the radio buttons is a button labeled 'Clear SMTP gateway settings'. At the bottom of the page is a text field labeled 'Custom 'From:' address' which contains 'labmanager.sg1@company.com'. On the left side of the screenshot, there are five numbered arrows pointing to specific elements: arrow 2 points to the 'New' button; arrow 3 points to the 'SMTP gateway name' radio button; arrow 4 points to the 'SMTP gateway IP' radio button; arrow 5 points to the 'Custom 'From:' address' text field.

2. Specify recipient e-mail addresses:
 - a. Click **New**; the Add List Item dialog displays.
 - b. Enter a recipient e-mail address.
 - c. Repeat for other recipients, if needed.
 - d. Click **OK** to close the dialog.
3. Specify mail sender information. The fields that are available depend on the version of SGOS 6.5.x you are using:
 - (SGOS 6.5.2 and later) In the **SMTP server** field, enter the SMTP server in one of the following formats:
 - Hostname of the server. The hostname can resolve to either an IPv4 or IPv6 address.
 - IPv4 or IPv6 address of the server.

Because the appliance does not validate the values you enter, make sure that they are correct before applying changes.

- (Versions earlier than SGOS 6.5.2) In the **SMTP gateway name** field, enter the host name of your mail server. The gateway name can resolve to either an IPv4 or IPv6 address.

In the **SMTP gateway IP** field, enter the IPv4 or IPv6 address of your mail server. The default port for SMTP is 25. If your configuration uses a different port, you can set the port number with the following CLI command:

```
SGOS#(config smtp) server {domain_name | IP_address} [port]
```

4. (Optional) The **Clear SMTP Settings** option clears the selected setting, but it does *not* delete the setting. For example, if you click **SMTP gateway name** and click **Clear SMTP Settings**, the value disappears. When you click **SMTP gateway name** again, the console displays the value.
5. (Optional) Specify the sender's email address in the **The "From:" field for the email** field. For example, enter the e-mail address of the lab manager responsible for administering ProxySG appliances.

Note: In earlier versions of SGOS 6.5.x, the field name is **Custom 'From' address**.

If you do not specify an email address here, event notifications display the name of the appliance for the sender's address. For information on configuring the appliance name, see ["Configuring the ProxySG Name"](#) on page 35.

6. Click **Apply**. The console confirms your changes.

Note: The email subject field states "ProxySG appliance" and is not configurable.

Syslog Event Monitoring

Syslog is an event-monitoring protocol that is especially popular in UNIX environments. Sites that use syslog typically have a log host node, which acts as a sink (repository) for several devices on the network. You must have a syslog daemon operating in your network to use syslog monitoring. The syslog format is: Date Time Hostname Event.

Most clients using syslog have multiple devices sending messages to a single syslog daemon. This allows viewing a single chronological event log of all of the devices assigned to the syslog daemon. An event on one network device might trigger an event on other network devices, which, on occasion, can point out faulty equipment.


If redundancy is necessary for your deployment, additional loghost servers can be configured for notification. When multiple loghosts are available, the event log message is sent simultaneously to multiple servers, reducing the possibility of data loss.

To retrieve event logs and view them on an external server, see "[Securely Retrieving Event Logs from the Appliance](#)" on page 1314.

Note: When a host is removed from the active syslog host list, a message indicating that syslog has been deactivated is sent to the host(s). This message alerts administrators that this host will no longer be receiving logs from this ProxySG.

To enable syslog monitoring:

1. Select the **Maintenance > Event Logging > Syslog** tab.



2. Click **New**. The **Add list item** displays.
3. In the **Add syslog loghost** field, enter the IPv4 or IPv6 address of your loghost server, or specify a domain name that resolves to an IPv4 or IPv6 address.
4. Click **OK**.
5. (Optional) Repeat steps 2-4 to add additional syslog servers to the loghost list.
6. Select **Enable Syslog**.
7. Click **Apply**.

Note: Event log messages are automatically sent to all syslog servers in the loghost list. Though the event log logs all events, the messages sent to the syslog server include only events matching the selected event logging levels. For more information, see ["Selecting Which Events to View"](#) on page 1309.

Related CLI Commands to Enable Syslog Monitoring

In addition to the Management Console, the CLI has more options for configuring Syslog monitoring:

```
SGOS#(config event-log) syslog add {hostname | IP_Address}
SGOS#(config event-log) syslog clear
SGOS#(config event-log) syslog {disable | enable}
SGOS#(config event-log) syslog facility {auth | daemon | kernel |
local0 | local1 | local2 | local3 | local4 | local5 | local6 | local7 |
lpr | mail | news | syslog | user | uucp}
SGOS#(config event-log) syslog remove {hostname | IP_Address}
```

Securely Retrieving Event Logs from the Appliance

As an alternative to sending logs to Syslog servers (see ["Syslog Event Monitoring"](#) on page 1312), you can retrieve event log data over a secure connection and save it on an external server.

To transfer event logs over a secure connection, the remote server periodically retrieves the event log data from the ProxySG appliance. With this in mind, you should configure event log settings on the appliance to make sure that the retrieved data reflects all current events.

Ensuring that Current Event Log Data is Retrieved

To ensure that you retrieve current event log data, Symantec recommends that you set the event log to overwrite older events. With this setting, the event log comprises two files, which are rotated when there is log overflow. For example, if the maximum size of 10 MB is reached and an event needs to be written to the log, the two log files are rotated, resulting in the loss of the older 5 MB file. To prevent data loss, you should configure event log settings so that the older 5 MB log file is retrieved before it is overwritten.

To determine an appropriate maximum size, you could consider the rate of event log growth, including factors such as logging levels and the typical number of events that occur in your environment. In turn, consider both log growth and the log's maximum size to determine how often the external server should retrieve log data. For assistance, refer to your Symantec Support Engineer.

To ensure that current event log data is retrieved:

1. Verify or change the event log's maximum size. See ["Setting Event Log Size"](#) on page 1310.
2. Specify that the event log should overwrite older events. See ["Setting Event Log Size"](#) on page 1310.

-
3. Set the frequency with which the external server retrieves event log data. Specify an amount of time that allows the server to retrieve all data before it is overwritten.

For example, consider an event log that logs all levels above Verbose and grows quickly. To ensure that no data is lost, you set the log size to a 10 MB maximum and specify that event log data is retrieved every 50 seconds. See the following example.

Example: Using cURL over HTTPS to Retrieve Event Log Data

Note: This example is not intended to be used for a real-world scenario; it could be inadequate for your purposes. This example is meant for demonstration purposes only.

The following is an example of using cURL to retrieve event log data over HTTPS every five seconds.

```
#!/bin/sh
# This script archives Event Log data into the specified file \
  (<archive_filename>).
# The script will collect new Event Log entries every 5 seconds by default.
# Specify a different time interval if you wish (<refresh-time-in-seconds>).
# You will need to set the username and userpass for your system.

username="admin"
userpass="admin"

if [ "$#" -lt 2 ]; then
    echo "usage: archive-eventlog <IP_Address> <archive_filename> \
      [refresh-time-in-seconds]"
    exit 1
fi

sg_addr=$1
saved=$2

# Optional refresh time
refresh_time=$3
if [ "${refresh_time}" = "" ]; then
    # Default is for 5 seconds
    refresh_time=5
fi

base_path="/tmp/"
temp_file="${base_path}tmp_eventlog.$$"
temp_file_ok="${base_path}tmp_eventlog_ok.$$"
rm -rf ${temp_file} ${temp_file_ok}

while true; do
```

```
curl https://${sg_addr}:8082/eventlog/fetch=0xffffffff -k --user \
${username}:${userpass} > ${temp_file} 2> /dev/null
if [ $? -eq 0 ]; then
    # We successfully got data downloaded
    # Pre-check that the download ends in a "good line" or fix it.
    eof_char=`tail -c 1 ${temp_file} | tr '\n' 'X'`
    if [ "${eof_char}" != "X" ]; then
        # Looks like the last line is incomplete
        # Let's trim away the incomplete line to clean up the data
        # Trimmed lines will be in the next refresh
        sed 'd' ${temp_file} > ${temp_file_ok}
    else
        mv ${temp_file} ${temp_file_ok}
    fi
    if [ -e ${saved} ]; then
        # We have previously archived data, so add to it
        last_line=`tail -1 ${saved}`
        # Test the last archived line is in the new content
        grep "${last_line}" ${temp_file_ok} 2>/dev/null 1>&2
        if [ $? -eq 0 ]; then
            # Add the content after the last matching line
            match_line=`echo "${last_line}" | tr / .`
            sed "0,/${match_line}/d" ${temp_file_ok} >> \
                ${saved} 2> /dev/null
        else
            # Nothing matched so add all data just downloaded
            cat ${temp_file_ok} >> ${saved} 2> /dev/null
        fi
    else
        # No previously archived data, so add all of the \
        data just downloaded
        cat ${temp_file_ok} > ${saved} 2> /dev/null
    fi
fi

# Cleanup before resting
rm -rf ${temp_file} ${temp_file_ok}
sleep ${refresh_time}

done
```

Viewing Event Log Configuration and Content

You can view the event log configuration, from `show` or from `view` in the event-log configuration mode.

To view the event log configuration:

At the prompt, enter the following command:

- ❑ From anywhere in the CLI

```
SGOS> show event-log configuration
Settings:
  Event level: severe + configuration + policy + informational
  Event log size: 10 megabytes
  If log reaches maximum size, overwrite earlier events
  Syslog loghost: <none>
  Syslog notification: disabled
  Syslog facility: daemon
Event recipients:
SMTP gateway:
  mail.heartbeat.bluecoat.com
```

-or-

❑ From the (config) prompt:

```
SGOS#(config) event-log
SGOS#(config event-log) view configuration
Settings:
  Event level: severe + configuration + policy + informational
  Event log size: 10 megabytes
  If log reaches maximum size, overwrite earlier events
  Syslog loghost: <none>
  Syslog notification: disabled
  Syslog facility: daemon
Event recipients:
SMTP gateway:
  mail.heartbeat.bluecoat.com
```

To view the event log contents:

You can view the event log contents from the `show` command or from the event-log configuration mode.

Note: The results displayed include events only for the configured event logging levels. For more information, see ["Selecting Which Events to View"](#) on page 1309.

The syntax for viewing the event log contents is

```
SGOS# show event-log
-or-
SGOS# (config event-log) view
[start [YYYY-mm-dd] [HH:MM:SS]] [end [YYYY-mm-dd] [HH:MM:SS]] [regex
regex | substring string]
```

Pressing <Enter> shows the entire event log without filters.

The order of the filters is unimportant. If `start` is omitted, the start of the recorded event log is used. If `end` is omitted, the end of the recorded event log is used.

If the date is omitted in either `start` or `end`, it must be omitted in the other one (that is, if you supply just times, you must supply just times for both `start` and `end`, and all times refer to today). The time is interpreted in the current time zone of the appliance.

Understanding the Time Filter

The entire event log can be displayed, or either a starting date/time or ending date/time can be specified. A date/time value is specified using the notation ([YYYY-MM-DD] [HH:MM:SS]). Parts of this string can be omitted as follows:

- ❑ If the date is omitted, today's date is used.
- ❑ If the time is omitted for the starting time, it is 00:00:00.
- ❑ If the time is omitted for the ending time, it is 23:59:59.

At least one of the date or the time must be provided. The date/time range is inclusive of events that occur at the start time as well as dates that occur at the end time.

Note: If the notation includes a space, such as between the start date and the start time, the argument in the CLI should be quoted.

Understanding the Regex and Substring Filters

A regular expression can be supplied, and only event log records that match the regular expression are considered for display. The regular expression is applied to the text of the event log record not including the date and time. It is case-sensitive and not anchored. You should quote the regular expression.

Since regular expressions can be difficult to write properly, you can use a substring filter instead to search the text of the event log record, not including the date and time. The search is case sensitive.

Regular expressions use the standard regular expression syntax as defined by policy. If both regex and substring are omitted, then all records are assumed to match.

Example

```
SGOS# show event-log start "2009-10-22 9:00:00" end "2009-10-22
9:15:00"
2009-10-22 09:00:02+00:00UTC "Snapshot sysinfo_stats has fetched /
sysinfo-stats " 0 2D0006:96 ../Snapshot_worker.cpp:183
2009-10-22 09:05:49+00:00UTC "NTP: Periodic query of server
ntp.bluecoat.com, system clock is 0 seconds 682 ms fast compared to NTP
time. Updated system clock. " 0 90000:1 ../ntp.cpp:631
```

Section D: Monitoring Network Devices (SNMP)

This section discusses the following topics:

- ❑ ["Introduction to SNMP"](#)
- ❑ ["About SNMP Traps and Informs" on page 1320](#)
- ❑ ["About Management Information Bases \(MIBs\)" on page 1322](#)
- ❑ ["Adding and Enabling an SNMP Service and SNMP Listeners" on page 1323](#)
- ❑ ["Configuring SNMP Communities" on page 1325](#)
- ❑ ["Configuring SNMP for SNMPv1 and SNMPv2c" on page 1326](#)
- ❑ ["Configuring SNMP for SNMPv3" on page 1330](#)

Introduction to SNMP

Simple Network Management Protocol (SNMP) is used in network management systems to monitor network devices for health or status conditions that require administrative attention. The ProxySG supports SNMPv1, SNMPv2c, and SNMPv3.

This section discusses the following topics:

- ❑ ["Typical Uses of SNMP"](#)
- ❑ ["Types of SNMP Management" on page 1319](#)
- ❑ ["Components of an SNMP Managed Network" on page 1320](#)

Typical Uses of SNMP

Some typical uses of SNMP include:

- ❑ Monitoring device uptimes
- ❑ Providing information about OS versions
- ❑ Collecting interface information
- ❑ Measuring network interface throughput

For more information, see the following sections.

Types of SNMP Management

The ProxySG provides the capability to configure SNMP for single network management systems, a multiple user NMS, and for notification only.

If you are not using a network manager to interrogate the state of the ProxySG, configure the ProxySG to provide required traps without any SNMP read-write operations. As a result, no ports are defined as listeners for SNMP. If any or all SNMP listeners in the services are deleted or disabled, you can still configure traps and informs to go out.

Components of an SNMP Managed Network

An SNMP managed network consists of the following:

- ❑ Managed devices—Network nodes that contain an SNMP agent and reside on a managed network.
- ❑ Agents—Software processes that respond to queries using SNMP to provide status and statistics about a network node.
- ❑ Network Management Systems (NMSs)—Each NMS consists of a combination of hardware and software used to monitor and administer a network. An NMS executes applications that monitor and control managed devices. You can have one or more NMSs on any managed network.

You can select the SNMP versions the ProxySG supports to match the configuration of your SNMP manager, as well as select the ports on which SNMP listens. SNMP traps and informs work over UDP only; SGOS does *not* support traps and informs over TCP connections, even if that is supported by your management tool.

You can configure the ProxySG to work with a sophisticated network environment with NMS users that have different access requirements for using SNMP than in a single NMS environment. For example, some users might have access to particular network components and not to others because of their areas of responsibility. Some users might have access based on gathering statistics, while others are interested in network operations.

See Also

- ❑ ["About Management Information Bases \(MIBs\)"](#)

About SNMP Traps and Informs

SNMP agents (software running on a network-connected device) not only listen for queries for data, but also can be configured to send traps or informs (alert messages) to a network-monitoring device that is configured to receive SNMP traps. The only difference between a trap and an inform is that the SNMP manager that receives an inform request acknowledges the message with an SNMP response; no response is sent for regular traps.

SNMP traps work with SNMPv1, SNMPv2c, and SNMPv3. SNMP informs work with SNMPv2c and SNMPv3 only.

When you have traps enabled, events that can trigger a trap to be sent include such things as hardware failures and elevations or decreases in component thresholds. The default SNMP traps and informs include the following standard SNMP traps:

- ❑ `coldStart`—signifies that the SNMP entity, supporting a notification originator application, is reinitializing itself and that its configuration might have been altered. This MIB is described in `SNMPv2-MIB.txt`.

-
- ❑ `warmStart`—The SNMP entity, supporting a notification originator application, is reinitializing itself such that its configuration is unaltered. This MIB is described in `SNMPv2-MIB.txt`.
 - ❑ `linkUp`—The SNMP entity, acting in an agent role, has detected that the `ifOperStatus` object for one of its communication links left the down state and transitioned into some other state (but not into the `notPresent` state). This other state is indicated by the included value of `ifOperStatus`. This MIB is described in `IF-MIB.txt`.
 - ❑ `linkDown`—The SNMP entity, acting in an agent role, has detected that the `ifOperStatus` object for one of its communication links is about to enter the downstate from some other state (but not from the `notPresent` state). This other state is indicated by the included value of `ifOperStatus`. This MIB is described in `IF-MIB.txt`.

The following traps require additional configuration:

- ❑ Authentication failure traps first must be enabled. See ["Configuring SNMP Communities"](#) on page 1325.
- ❑ The attack trap occurs if attack detection is set up. See ["Preventing Denial of Service Attacks"](#) on page 1279.
- ❑ The disk/sensor traps are driven by the health monitoring settings (as is the health monitoring trap). See ["Changing Threshold and Notification Properties"](#) on page 1347.
- ❑ The health check trap occurs if it is set up in the health check configuration. See ["Configuring Health Check Notifications"](#) on page 1369.
- ❑ The policy trap goes off if there is policy to trigger it. Refer to the *Visual Policy Manager Reference* or the *Content Policy Language Guide*. Many of the feature descriptions throughout this guide also include information about setting policy.

See Also

- ❑ ["Configuring SNMP Communities"](#) on page 1325
- ❑ ["Changing Threshold and Notification Properties"](#) on page 1347
- ❑ ["Adding Community Strings for SNMPv1 and SNMPv2c"](#)
- ❑ ["Configuring SNMP Traps for SNMPv1 and SNMPv2c"](#)
- ❑ ["Configuring SNMP for SNMPv3"](#)
- ❑ ["Configuring SNMP Traps and Informs for SNMPv3"](#)

About Management Information Bases (MIBs)

A Management Information Base (MIB) is a text file (written in the ASN.1 data description language) that contains the description of a managed object. SNMP uses a specified set of commands and queries, and the MIBs contain information about these commands and the target objects.

One of the many uses for MIBs is to monitor system variables to ensure that the system is performing adequately. For example, a specific MIB can monitor variables such as temperatures and voltages for system components and send traps when something goes above or below a set threshold.

The Symantec MIB specifications adhere to RFC1155 (v1-SMI), RFC1902 (v2-SMI), RFC1903 (v2-TC), and RFC1904 (v2-CONF.)

Note: Some common MIB types, such as 64-bit counters, are not supported by SNMPv1. We recommend using either SNMPv2c or, for best security, SNMPv3.

The ProxySG uses both public MIBs and Symantec proprietary MIBs. You can download the MIB files from MySymantec.

To download the MIBs:

1. Go to MySymantec:
<https://support.symantec.com>
2. Select **Downloads > Network Protection (Blue Coat) Downloads**.
3. When prompted, log in with your MySymantec credentials.
4. Select your product.
5. Select your appliance model (if applicable).
6. Select a software version.
7. Accept the License Agreement.
8. Select the file(s) to download and click **Download Selected Files**.

Note: The first time you download files, you are prompted to install the Download Manager. Follow the onscreen prompts to download and run the installer. For more information, refer to <https://www.symantec.com/support-center/getting-started>.

9. The Download Manager window opens. Select the download location.


Note: Complete instructions are also available online at:
<https://www.symantec.com/support-center/getting-started>
Bookmark this page for future reference.

Adding and Enabling an SNMP Service and SNMP Listeners


There is one disabled SNMP listener defined by default on the ProxySG, which you can delete or enable, as needed. You can also add additional SNMP services and listeners. Although you can configure traps and informs to go out if all the SNMP listeners are deleted or disabled, configuring SNMP listeners sets up the UDP ports the ProxySG uses to listen for SNMP commands. The service ports set up for *listening* to SNMP requests are independent of the trap or inform addresses and ports specified for *sending* traps.

To add and enable an SNMP service and listeners:

1. Select the **Configuration > Services > Management Services** tab.
2. Click **Add**. The New Service dialog displays.



3. Enter a name for the SNMP Service.
4. In the **Services** drop-down list, select **SNMP**.
5. Click **New**. The New Listener dialog displays.



6. Configure listener options:
 - a. In the **Destination addresses** area, select **All SG IP addresses** or select **IP Address** and select a specific IP address from the drop-down list. The IP address can be either IPv4 or IPv6.
 - b. Enter the port for this listener.
 - c. Select **Enabled** to enable this listener.
 - d. Click **OK** to close the New Listener dialog, then click **OK** again to close the New Service dialog.
7. Click **Apply**.

To delete an SNMP service:

1. Select **Configuration > Services > Management Services**. The Management Services tab displays.
2. Select the SNMP service to delete and click **Delete**. A dialog box prompts you to confirm the deletion.
3. Click **OK** to delete the SNMP service, then click **Apply**.

To delete an SNMP listener:

1. Select the **Configuration > Services > Management Services** tab.
2. Select an SNMP service in the list and click **Edit**. The Edit Service dialog displays.
3. Select the listener to delete and click **Delete**. A dialog box prompts you for confirmation.
4. Click **OK** to delete the listener, then click **OK** again to close the Edit Service dialog.
5. Click **Apply**.

See Also


- ❑ ["Managing Proxy Services" on page 109](#)

Configuring SNMP Communities

For the ProxySG to listen for SNMP commands, you must enable at least one SNMP listener. After you add and enable an SNMP service (see ["Adding and Enabling an SNMP Service and SNMP Listeners"](#) on page 1323), you are ready to configure SNMP communities and users and enable traps and informs (see ["About SNMP Traps and Informs"](#) on page 1320).

To configure SNMP:

1. Select the **Maintenance > SNMP > SNMP General** tab.




2. In the **Protocols** area, **SNMPv1**, **SNMPv2**, and **SNMPv3** are all enabled by default. Select the specific versions that match the configuration of your SNMP manager.

Note: Only **SNMPv3** uses the Engine ID, which is required to be unique among SNMP agents and systems that are expected to work together.

The Engine ID is set by default to a value that is derived from the ProxySG serial number and the Symantec SNMP enterprise code. This is a unique hexadecimal identifier that is associated with the ProxySG. It appears in each SNMP packet to identify the source of the packet. The configured bytes must *not* all be equal to zero or to 0FFH (255).

If you reset the engine ID and want to return it to the default, click **Set to Default**. You do not need to reboot the system after making configuration changes to SNMP.




3. In the **Traps and Informs** area, enable traps and informs, as required.
 - a. Select **Enable use of traps and informs** to enable SNMP traps (for SNMPv1, SNMPv2c, and SNMPv3) or informs (for SNMPv2c and SNMPv3 only).

- b. Select **Enable SNMP authentication failure traps** to have an SNMP authentication failure trap sent when the SNMP protocol has an authentication failure.

Note: For SNMPv1 and SNMPv2c, this happens when the community string in the SNMP packet is not correct (does not match one that is supported). For SNMPv3, this happens when the authentication hash of an SNMP packet is not correct for the specified user.

- c. To perform a test trap, click **Perform test trap**, enter the trap data (string) to be sent, and click **Execute Trap**. This sends a policy notification, as defined in the BLUECOAT-SG-POLICY-MIB, to all configured trap and inform recipients, and it is intended as a communications test.



4. In the **sysContact** field, enter a string that identifies the person responsible for administering the appliance.
5. In the **sysLocation** field, enter a string that describes the physical location of the appliance.
6. Click **Apply**.

See Also

- ❑ ["Monitoring Network Devices \(SNMP\)"](#)
- ❑ ["Adding and Enabling an SNMP Service and SNMP Listeners"](#)
- ❑ ["Adding Community Strings for SNMPv1 and SNMPv2c"](#)
- ❑ ["Configuring SNMP Traps for SNMPv1 and SNMPv2c"](#)
- ❑ ["Configuring SNMP for SNMPv3"](#)
- ❑ ["Configuring SNMP Traps and Informs for SNMPv3"](#)

Configuring SNMP for SNMPv1 and SNMPv2c


Community strings are used for SNMPv1 and SNMPv2c only. SNMPv3 replaces the use of a community string with the ability to define a set of users. See ["Configuring SNMP for SNMPv3"](#) on page 1330.

Adding Community Strings for SNMPv1 and SNMPv2c


Community strings restrict access to SNMP data. After you define a community string, you set an authorization mode of either *read* or *read-write* to allow access using that community string. The mode *none* allows you to use a community string for traps and informs only.

To add a community string:

1. Select the **Maintenance > SNMP > SNMPv1-v2c Communities** tab.



2. Click **New**. The Create Community dialog displays.



3. In the **Community String** field, name the new string.
4. In the **Authorization** field, select the authorization level (**None**, **Read-only**, or **Read-write**).

5. To use all available source addresses, click **OK** and proceed to Step 7.
6. To configure an access control list (available if you selected **Read-only** or **Read-write**), select **Enforce access control list for requests** and click **Edit ACL**. The Source Addresses dialog displays.
 - a. Click **Add**. The Add IP/Subnet dialog displays.
 - b. Enter the IP/Subnet Prefix and the Subnet Mask, then click **OK** in all open dialogs until you return to the **SNMPv1-v2c Communities** tab.
7. Click **Apply**.

To edit a community string:

1. Select the **Maintenance > SNMP > SNMPv1-v2c Communities** tab.
2. Select the community string to edit and click **Edit**. The Edit (*community name*) dialog displays.
3. Edit the parameters as required, then click **OK**.
4. Click **Apply**.

See Also

- ❑ ["Adding and Enabling an SNMP Service and SNMP Listeners"](#)
- ❑ ["Configuring SNMP for SNMPv1 and SNMPv2c"](#)
- ❑ ["Configuring SNMP Users for SNMPv3"](#)
- ❑ ["Monitoring Network Devices \(SNMP\)"](#)


Configuring SNMP Traps for SNMPv1 and SNMPv2c

The ProxySG can send SNMP traps (for SNMPv1 and SNMP v2c) and informs (for SNMPv2c) to a management station as they occur. Each SNMP notification is sent to all defined trap and inform receivers (of all protocols). You can also enable authorization traps to send notification of attempts to access the Management Console.

If the system reboots for any reason, a *cold start trap* is sent. A *warm start trap* is sent if a you perform a software-only reboot without a hardware reset. No configuration is required.

To add SNMP traps:

1. Select the **Maintenance > SNMP > SNMP v1-v2c Traps** tab.
2. Click **New**. The Create Trap or Inform Receiver dialog displays.



3. From the **Community string** drop-down list, select a previously created community string (see ["Configuring SNMP Communities"](#) on page 1325)
4. Select the **Type** of trap. The difference between a trap and an inform is that the SNMP manager that receives an inform request acknowledges the message with an SNMP response. No response is sent for regular traps.
5. In the **Receiver** area, enter the IP address and port number.
6. Click **OK**, then click **Apply**.

To edit a trap or inform:

1. Select the **Maintenance > SNMP > SNMP v1-v2c Traps** tab.
2. Select a trap or inform in the list and click **Edit**. The Edit (*trap name*) Trap or Inform Receiver dialog displays.
3. Edit the settings as desired and click **OK**.
4. Click **Apply**.

See Also

- ❑ ["Monitoring Network Devices \(SNMP\)"](#)
- ❑ ["About SNMP Traps and Informs"](#)
- ❑ ["Adding and Enabling an SNMP Service and SNMP Listeners"](#)

- ❑ ["Configuring SNMP Communities"](#)
- ❑ ["Adding Community Strings for SNMPv1 and SNMPv2c"](#)
- ❑ ["Configuring SNMP for SNMPv3"](#)
- ❑ ["Configuring SNMP Users for SNMPv3"](#)
- ❑ ["Configuring SNMP Traps and Informs for SNMPv3"](#)

Configuring SNMP for SNMPv3

For SNMPv v3, you configure users instead of community strings. You then configure the traps and informs by user rather than by community string.

This section discusses the following topics:

- ❑ ["About Passphrases and Localized Keys"](#)
- ❑ ["Configuring SNMP Users for SNMPv3"](#) on page 1330
- ❑ ["Configuring SNMP Traps and Informs for SNMPv3"](#) on page 1333

About Passphrases and Localized Keys

Although it is optional to use passphrases or localized keys, using one or the other provides the increased security of SNMPv3. For most deployments, passphrases provide adequate security. For environments in which there are increased security concerns, you have the option of setting localized keys instead of passphrases. In the configuration, if you set a passphrase, any localized keys are immediately deleted and only the passphrase remains. If you set a localized key, any passphrase is deleted and the localized key is used.

If you need to use localized keys, you can enter one for the ProxySG and add keys for other specified Engine IDs. Since the ProxySG acts as an agent, its localized key is all that is needed to conduct all SNMP communications, with the single exception of SNMP informs. For informs, you need to provide the localized key that corresponds to each engine ID that is going to receive your informs.

Configuring SNMP Users for SNMPv3

When you set up users, you configure authentication and privacy settings, as required.

Note: The enhanced security of SNMPv3 is based on each user having an authentication passphrase and a privacy passphrase. For environments in which there are increased security concerns, you have the option of setting up localized keys instead of passphrases.

You can enable authentication without enabling privacy, however, you cannot enable privacy without enabling authentication. In an authentication-only scenario, a secure hash is done so the protocol can validate the integrity of the packet. Privacy adds the encryption of the packet data.

To configure SNMP users:

1. Select the **Maintenance > SNMP > SNMPv3 Users** tab.
2. Click **New**. The Create User dialog displays.

The screenshot shows the 'Create User' dialog box with the following fields and options:

- User name:** A text field containing 'boyson' (labeled 3).
- Authentication section:**
 - Mode:** Radio buttons for 'None', 'MD5' (selected), and 'SHA' (labeled 4a).
 - Buttons:** 'Change Passphrase' (labeled 4b) and 'Set Localized Keys' (labeled 4c).
- Privacy section:**
 - Mode:** Radio buttons for 'None', 'DES' (selected), and 'AES' (labeled 5a).
 - Buttons:** 'Change Passphrase' (labeled 5b) and 'Set Localized Keys' (labeled 5c).
- Authorization section:**
 - Mode:** Radio buttons for 'None', 'Read-only', and 'Read-write' (selected) (labeled 6).
- Buttons:** 'OK' and 'Cancel' at the bottom.

3. Enter the name of the user.
4. In the **Authentication** area:
 - a. Select the authentication mode: **MD5** (Message Digest Version 5) or **SHA** (Secure Hash Algorithm).
 - b. Click **Change Passphrase** to set or change the authentication passphrase. If your environment requires a higher level of security, you have the option of setting up localized keys instead of passphrases. See Step c. Enter and confirm the passphrase, then click **OK**.
 - c. (Optional) To set up localized keys for authentication instead of using an authentication passphrase, click **Set Localized Keys**. The Localized Keys dialog displays. When you set up localized keys, any password is deleted and the localized keys are used instead.
 - Click **New**. The Set Localized Key dialog displays.
 - If the Engine ID is Self, enter and confirm the localized key (hexadecimal), then click **OK**.

- To add additional localized keys, enter the Engine ID (hexadecimal) and the localized key, then click **OK**.
5. In the **Privacy** area:
 - a. To set up the privacy mode, select **DES** (Data Encryption Standard) or **AES** (Advanced Encryption Standard).
 - b. Click **Change Passphrase** to set or change the privacy passphrase. If your environment requires a higher level of security, you have the option of setting up localized keys instead of passphrases. See Step c.
 - Enter and confirm the passphrase, then click **OK**.
 - c. (Optional) To set up localized keys for privacy instead of using a privacy passphrase, click **Set Localized Keys**. The Localized Keys dialog displays. If you have set up a privacy passphrase, you will not be able to set up localized keys.
 - Click **New**. The Set Localized Key dialog displays.
 - If the Engine ID is Self, enter and confirm the localized key (hexadecimal), then click **OK**.
 - To add additional localized keys, enter the Engine ID (hexadecimal) and the localized key, then click **OK**.
 6. Select the **Authorization** mode for this user: **None**, **Read-only**, or **Read-write**.
 7. Click **OK** to close the Create User dialog.
 8. Click **Apply**.

To edit a user:

1. Select **Maintenance > SNMP > SNMPv3 Users**.
2. Select the user to edit and click **Edit**. The Edit (*user name*) dialog displays.
3. Edit the parameters as required, then click **OK**.
4. Click **Apply**.

For a complete list of the CLI commands to edit an SNMPv3 user, see Chapter 3 “Privileged Mode Commands” in the *Command Line Interface Reference*.

See Also


- ❑ ["Configuring SNMP Communities"](#)
- ❑ ["Adding Community Strings for SNMPv1 and SNMPv2c"](#)
- ❑ ["Configuring SNMP Traps for SNMPv1 and SNMPv2c"](#)
- ❑ ["Configuring SNMP Traps and Informs for SNMPv3"](#)

Configuring SNMP Traps and Informs for SNMPv3

Before you can configure SNMPv3 traps and informs, you must set up users and their associated access control settings. (See "Configuring SNMP for SNMPv3" on page 1330.) The difference between a trap and an inform is that the SNMP manager that receives an inform request acknowledges the message with an SNMP response; no response is sent for regular traps.

To configure SNMP traps for SNMPv3:

1. Select the **Maintenance > SNMP > SNMPv3 Traps** tab.
2. Click **New**. The Create Trap or Inform Receiver dialog displays.



3. Select the user from the drop-down list.
4. Select **SNMPv3 Trap** or **SNMPv3 Inform**.
5. In the **Receiver** section, enter the IP address and port number.
6. Click **OK**, then click **Apply**.

To edit a trap or inform:

1. Select the **Maintenance > SNMP > SNMPv3 Traps** tab.
2. Select a trap or inform in the list and click **Edit**. The Edit (*trap name*) Trap or Inform Receiver dialog displays.
3. Edit the settings as desired and click **OK**.
4. Click **Apply**.

For the full list of subcommands to edit traps and informs for SNMPv3 users, see Chapter 3 "Privileged Mode Commands" in the *Command Line Interface Reference*.

See Also

- ❑ ["About SNMP Traps and Informs"](#)
- ❑ ["Configuring SNMP Communities"](#)
- ❑ ["Adding Community Strings for SNMPv1 and SNMPv2c"](#)
- ❑ ["Configuring SNMP Traps for SNMPv1 and SNMPv2c"](#)

Section E: Configuring Health Monitoring

The health monitor records the aggregate health of the ProxySG, by tracking status information and statistics for select resources, and aids in focusing attention when a change in health state occurs. On the ProxySG, the health monitor tracks the status of key hardware components (such as the thermal sensors, and CPU use), and the health status for configured services (such as ADN). When the health monitor detects deviations in the normal operating conditions of the device, the health status changes.

Note: The change in health status is displayed in the Management Console and by the status LED on the appliance.

A change in health status does not always indicate a problem that requires corrective action; it indicates that a monitored metric has deviated from the normal operating parameters. The health monitor aids in focusing attention to the possible cause(s) for the change in health status.

In [Figure 72–1](#) below, the **Health:** monitor displays the overall health of the ProxySG in one of three states, **OK**, **Warning**, or **Critical**. Click the link to view the **Statistics > Health Monitoring** page, which lists the status of the system's health monitoring metrics.




Figure 72–1 Health Monitor as displayed on the Management Console

See Also

- ["About Health Monitoring"](#)

About Health Monitoring

Health Monitoring allows you to set notification thresholds on various internal metrics that track the health of a monitored system or device. Each metric has a *value* and a *state*.

The *value* is obtained by periodically measuring the monitored system or device. In some cases, the value is a percentage or a temperature measurement; in other cases, it is a status such as `Disk Present` or `Awaiting Approval`.

The *state* indicates the condition of the monitored system or device:

- ❑ **OK**—The monitored system or device is behaving within normal operating parameters.
- ❑ **WARNING**—The monitored system or device is outside typical operating parameters and may require attention.
- ❑ **CRITICAL**—The monitored system or device is failing, or is far outside normal parameters, and requires immediate attention.

The current state of a metric is determined by the relationship between the value and its monitoring *thresholds*. The Warning and Critical states have thresholds, and each threshold has a corresponding *interval*.

All metrics begin in the **OK** state. If the value crosses the Warning threshold and remains there for the threshold's specified interval, the metric transitions to the Warning state. Similarly, if the Critical threshold is exceeded for the specified interval, the metric transitions to the Critical state. Later (for example, if the problem is resolved), the value drops back down below the Warning threshold. If the value stays below the Warning threshold longer than the specified interval, the state returns to OK.

Every time the state changes, a notification occurs. If the value fluctuates above and below a threshold, no state change occurs until the value stays above or below the threshold for the specified interval of time.

This behavior helps to ensure that unwarranted notifications are avoided when values vary widely without having any definite trend. You can experiment with the thresholds and intervals until you are comfortable with the sensitivity of the notification settings.

Health Monitoring Example

Figure 72-2 provides an example of health monitoring. The graph is divided into horizontal bands associated with each of the three possible states. The lower horizontal line represents the Warning threshold and the upper horizontal line is the Critical threshold. The vertical bands represent 5 second time intervals.

Assume both thresholds have intervals of 20 seconds, and that the metric is currently in the OK state.

1. At time 0, the monitored value crosses the Warning threshold. No transition occurs yet. Later, at time 10, it crosses the critical threshold. Still, no state change occurs, because the threshold interval has not elapsed.
2. At time 20, the value has been above the warning threshold for 20 seconds--the specified interval. The state of the metric now changes to Warning, and a notification is sent. Note that even though the metric is currently in the critical range, the State is still Warning, because the value has not exceeded the Critical threshold long enough to trigger a transition to Critical.
3. At time 25, the value drops below the Critical threshold, having been above it for only 15 seconds. The state remains at Warning.

4. At time 30, it drops below the Warning threshold. Again the state does not change. If the value remains below the warning threshold until time 50, then the state will change to OK.
5. At time 50, the state transitions to OK. This transition occurs because the monitored value has remained below the Warning threshold for the configured interval of 20 seconds.




Figure 72–2 Relationship between the threshold value and threshold interval

Health Monitoring Cycle

The health monitoring process is a cycle that begins with the health state at OK. When the health monitor detects a change in the value of a monitored metric, the health state changes. The **Health:** indicator reflects the change in status.

Note: A change in health status does not always indicate a problem that requires corrective action; it indicates that a monitored metric has deviated from the normal operating parameters.

The **Health:** indicator is always visible in the Management Console, and the color and text reflect the most severe health state for all metrics— red for **Critical**, yellow for **Warning**, and green for **OK**. In the **Health Monitoring > Statistics** panel, the tabs for **General**, **License**, and **Status** metrics change color to reflect the most severe state of the metrics they contain. You might click the tabs to view the problem and assess the information. Based on the cause for the alert, the administrator might take diagnostic action or redefine the *normal* operating parameters for the metric and restore the health state of the ProxySG.

For example, if the revolutions per minute for **Fan 1 Speed** falls below the warning threshold, the appliance's health transitions to **Warning**. Because **Fan 1 Speed** is a metric in the **Status** tab, the **Statistics > Health Monitoring > Status** tab turns yellow. By clicking the **Health:** link and navigating to the yellow tab, you can view the alert. You might then examine the fan to determine whether it needs to be replaced (due to wear and tear) or if something is obstructing its movement.

To facilitate prompt attention for a change in the health state, you can configure notifications on the appliance.

Planning Considerations for Using Health Monitoring

The health monitor indicates whether the ProxySG is operating within the default parameters set on the appliance. Symantec recommends that you review these settings and adjust them to reflect the normal operating parameters for your environment. You can configure:

- ❑ Thresholds, to define what measurements generate warnings or critical alerts. See ["Changing Threshold and Notification Properties"](#) on page 1347.
- ❑ Time intervals, that determine whether a threshold has been crossed and whether an alert should be sent. See ["Changing Threshold and Notification Properties"](#) on page 1347.
- ❑ The means by which alerts are delivered, any combination of e-mail, SNMP trap, event log, or none. See [Section C: "Configuring Event Logging and Notification"](#) on page 1309 for more information.

About the Health Monitoring Metric Types

The ProxySG monitors the status of the following metrics:

- ❑ Hardware — Disk, Voltage, Temperature, Fan speed, Power supply
- ❑ System Resources — CPU, Memory, and Network usage
- ❑ ADN Status
- ❑ License Expiration and Utilization
- ❑ Cloud Services: Common Policy Communication Status
- ❑ ICAP—Number of queued and deferred connections
- ❑ Health Check Status — Health status of the services used by the appliance

- ❑ Subscription Communication Status — Database downloads for Application Protection, CachePulse, and Geolocation

These health monitoring metrics are grouped as General, Licensing, or Status metrics.

The system resources and licensing thresholds are user-configurable, meaning that you can specify the threshold levels that will trigger an alert.

The hardware and ADN status metrics are *not* configurable and are preset to optimal values. For example, on some platforms, a Warning is triggered when the CPU temperature reaches 55 degrees Celsius.

The health check status metric is also *not* configurable. It takes into account the most acute value amongst the configured health checks and the *severity* component for each health check.

Severity of a health check indicates how the value of a failed health check affects the overall health of the ProxySG, as indicated by the health monitor.

If, for example, three health checks are configured on the ProxySG:

- ❑ `dns.192.0.2.4` with severity *No-effect*
- ❑ `fwd.test` with severity *Warning*
- ❑ `auth.service` with severity *Critical*

The value of the health check status metric adjusts in accordance with the success or failure of each health check and its configured severity as shown below:

If all three health checks report healthy, the health check status metric is OK.

If `dns.192.0.2.4` reports unhealthy, the health check status remains OK. The health check status metric does not change because its severity is set to no-effect.

If `fwd.test` reports unhealthy, the health check status transitions to Warning. This transition occurs because the severity for this health check is set to warning.

If `auth.service` reports unhealthy, the health check status becomes Critical because its severity is set to critical.

Subsequently, even if `fwd.test` reports healthy, the health check status remains critical as `auth.service` reports unhealthy.

The health check status transitions to OK only if both `fwd.test` and `auth.service` report healthy.

Table 72–1 Health Check Status Metric — Combines the Health Check Result and the Severity Option

Configured Health Checks	Reporting as...						
<code>dns.192.0.2.4</code> severity: no-effect	Healthy	Unhealthy	Unhealthy	Healthy	Healthy	Healthy	Healthy
<code>fwd.test</code> severity: warning	Healthy	Healthy	Unhealthy	Unhealthy	Unhealthy	Healthy	Healthy

auth.service severity: critical	Healthy	Healthy	Healthy	Healthy	Unhealthy	Unhealthy	Healthy
Health Check Status	OK	OK	Warning	Warning	Critical	Critical	OK

You can configure the default **Severity** for all health checks in the **Configuration > Health Checks > General > Default Notifications** tab. For more information on configuring the severity option for health checks, see [Chapter 73: "Verifying the Health of Services Configured on the ProxySG"](#) on page 1353.

About the General Metrics

The following table lists the metrics displayed in the **Maintenance > Health Monitoring > General** page. The thresholds and intervals for these metrics are user-configurable.

To view the statistics on CPU utilization and memory utilization on the ProxySG, see ["Viewing System Statistics"](#) on page 682.

To view the statistics on interface utilization, see ["Viewing Efficiency and Performance Metrics"](#) on page 26.

Table 72–2 General Health Monitoring Metrics

Metric	Default Values		Notes
	Critical Threshold / Interval	Warning Threshold / Interval	
CPU Utilization	95% / 120 seconds	80% / 120 seconds	Measures the value of the primary CPU on multi-processor systems — <i>not</i> the average of all CPU activity.
Memory Utilization	95% / 120 seconds	90% / 120 seconds	Measures memory use and tracks when memory resources become limited, causing new connections to be delayed.
Interface Utilization	90% / 120 seconds	60% / 120 seconds	Measures the traffic (in and out) on the interface to determine if it is approaching the maximum capacity. (bandwidth maximum)
Cloud Services: Common Policy Communication Status	48 hours	24 hours	Monitors the success of cloud common policy synchronization. If a sync fails for 24 hours, a warning is issued.

Table 72–2 General Health Monitoring Metrics (Continued)

ICAP Connections	80%/120 seconds	N/A	Sets alert notifications for queued and deferred ICAP connections.

See Also:

- ❑ ["Changing Threshold and Notification Properties"](#) on page 1347
- ❑ ["Snapshot of the Default Threshold Values and States"](#) on page 1345
- ❑ ["Health Monitoring Cycle"](#) on page 1337
- ❑ ["Health Monitoring Example"](#) on page 1336

About the Licensing Metrics

Table 72–3 lists the metrics displayed in the **Maintenance > Health Monitoring > Licensing** page. On the license page, you can monitor the utilization of user-limited licenses and the expiration of time-limited licenses. Licenses that do not expire or do not have a user limit are not displayed because there is no need to monitor them for a change in state that could affect the ProxySG appliance's health.

The threshold values for license expiration metrics are set in days until expiration. In this context, a critical threshold indicates that license expiration is imminent. Thus, the Critical threshold value should be smaller than the Warning threshold value. For example, if you set the Warning threshold to 45, an alert is sent when there are 45 days remaining in the license period. The Critical threshold would be less than 45 days, for example 5 days.

For license expiration metrics, the threshold interval is irrelevant and is set to 0.

Note: For new ProxySG appliances running SGOS 5.3 or higher, the default Warning threshold for license expiration is 15 days.

For ProxySG appliances upgrading from earlier versions to SGOS 5.4, the default Warning threshold remains at the same value prior to the upgrade. For example, if the Warning threshold was 30 days prior to the upgrade, the Warning threshold will remain at 30 days after the upgrade.

Refer to the most current *Release Notes* for SGOS upgrade information.

Table 72–3 Licensing Health Monitoring Metrics

Metric	Default Values		Notes
	Critical Threshold / Interval	Warning Threshold / Interval	
User License Utilization	90% / 120 seconds	80% / 120 seconds	Monitors the number of users using the ProxySG.
SGOS Base and SSL Proxy License Expiration	0 days / 0	15 days / 0 (For new ProxySG appliances running SGOS 5.3 or higher; see note below)	Warns of impending license expiration.
		30 days / 0 (For non-new ProxySG appliances upgrading from earlier versions of SGOS)	
Cloud Services: Common Policy Expiration	0 days / 0	30 days / 0	Monitors the days until entitlement expiration.
Geolocation Expiration/ CachePulse Expiration/ Application Protection Expiration	0 days/0	30 days / 0	Monitors the days until entitlement expiration.
License Server Communication Status	0 days/0	6 days / 0	Monitors the appliance's ability to connect to the license validation server.
License Validation Status	0 days/0	30 days / 0	Detects license validity.
Subscription Communication Status	20 days/0	10 days/0	Detects download errors for the following features (if you have a valid and enabled subscription): • Application Protection • CachePulse • Geolocation

See Also

- ["About User Limits"](#) on page 136.

- ❑ "Tasks for Managing User Limits" on page 138.
- ❑ Chapter 3: "Licensing" on page 43.

About the Status Metrics

The following table lists the metrics displayed in the **Maintenance > Health Monitoring > Status** page. The thresholds for these metrics are *not* user-configurable.

Table 72–4 Status Health Monitoring Metrics

Metric	Threshold States and Corresponding Values
Disk Status	Critical: Bad Warning: Removed Offline Present (failing) Present (unsupported failing) OK: Not Present Present Present (unsupported)
Temperature — Motherboard and CPU	Threshold states and values vary by ProxySG models
Fan Speed	Threshold states and values vary by ProxySG models
Voltage — Bus Voltage, CPU Voltage, Power Supply Voltage	Threshold states and values vary by ProxySG models

Table 72–4 Status Health Monitoring Metrics (Continued)

ADN Connection Status	OK: Connected Connecting Connection Approved Disabled Not Operational Warning: Approval Pending Mismatching Approval Status Partially Connected Critical: Disconnected Connection Denied See <i>"Reviewing ADN Health Metrics"</i> on page 762 for more information about the ADN metrics.
ADN Manager Status	OK: Not a Manager No Approvals Pending Warning: Approvals Pending
Health Check Status	OK: No health checks with <i>Severity: Warning</i> or <i>Critical</i> are failing. A health check with <i>Severity: No-effect</i> might be failing. Warning: One or more health checks with <i>Severity: Warning</i> has failed. Critical: One or more health checks with <i>Severity: Critical</i> has failed.

Table 72–4 Status Health Monitoring Metrics (Continued)

Subscription Communication Status	OK: No update errors Warning: 10 or more subsequent database downloads (after the first successful one) have failed. Critical: - When the feature is first enabled, the initial attempt to download the database failed. Error(s): CachePulse failed on initial download Geolocation failed on initial download Application Protection failed on initial download 20 or more subsequent database downloads (after the first successful one) have failed.
Sensor Count Status	On platforms that support it, this metric indicates if environmental sensors (which monitor temperature, fan speed, and voltage) are operational when the appliance boots up.
Reboot	Informational only: warm restart System rebooted with the <code>restart regular</code> command. cold restart System rebooted with <code>restart upgrade</code> command (Maintenance > Upgrade > Restart) or non-user initiated reboot, for example, power loss.
Failover	Informational only. If a failover occurs, notification is sent by the new master: <i>yyyy-mm-dd hh:mm:ss timezone:</i> <i>master_device_identifier</i> failed. <i>Appliance_name</i> is the new master.

Snapshot of the Default Threshold Values and States

See the table below for a quick glance at the health states and their corresponding threshold values:

Table 72–5 Health States and the Default Values for the Health Monitoring Metrics

General	Health States and Corresponding Default Values		
Metric	OK	Warning	Critical
CPU Utilization	less than 80%	80%	95%
Memory Utilization	less than 90%	90%	95%
Interface Utilization	less than 60%	60%	90%
Cloud Services: Common Policy Error Status	less than 24 hours since last successful update	24 hours	48 hours

Licensing	States and Corresponding Values		
Metric	OK	Warning	Critical
License Utilization	less than 80%	80%	90%
License Expiration	more than 15 days* more than 30 days **	15 days* 30 days**	0 days 0 days
Cloud Services: Common Policy Expiration	more than 30 days	30 days	0 days
Application Protection Expiration CachePulse Expiration Geolocation Expiration	more than 30 days	30 days	0 days
License Server Communication Status	more than 6 days	6 days	0 days
License Validation Status	more than 30 days	30 days	0 days

*For new ProxySG appliances running SGOS 5.3 or higher

** For non-new ProxySG appliances upgrading from earlier versions of the SGOS

Status	States and Corresponding Values		
Metric	OK	Warning	Critical
Disk status	Present/Not Present/ Present (unsupported)	Removed/Present (failing)/Present (unsupported failing)	Error
Temperature	Vary by ProxySG models		
Fan Speed	Vary by ProxySG models		
Voltage	Vary by ProxySG models		
ADN Connection Status	Connected Connecting Connection Approved Disabled Not Operational	Approval Pending Mismatching Approval Status Partially Connected	Disconnected Connection Denied
ADN Manager Status	Not a Manager No Approvals Pending	Approval Pending	
Health Check Status	No health checks with <i>Severity: Warning</i> or <i>Critical</i> are failing. A health check with <i>Severity: No-effect</i> might be failing.	One or more health checks with <i>Severity: Warning</i> has failed.	One or more health checks with <i>Severity: Critical</i> has failed.
Subscription Communication Status	No update errors	Numerous successive download errors	Before a successful initial download: <feature> failed on initial download After a successful initial download: Excessive successive download errors

Changing Threshold and Notification Properties

You can change the thresholds for the metrics in the **General** and **Licensing** tab to suit your network requirements. For the defaults, see "[About the Health Monitoring Metric Types](#)" on page 1338 and "[Viewing Health Monitoring Statistics](#)" on page 1349 for more information.


For health monitoring notifications, by default, all alerts are written to the event log. Any combination of the following types of notification can be set:

- ❑ **Log:** Inserts an entry into the Event log. See [Section C: "Configuring Event Logging and Notification"](#) on page 1309 for more information.
- ❑ **SNMP trap:** Sends an SNMP trap to all configured management stations. See [Section D: "Monitoring Network Devices \(SNMP\)"](#) on page 1319 for more information
- ❑ **E-mail:** Sends e-mail to all persons listed in the Event log properties. To use this option, you must add the recipient list to the Event log **Mail** option and ensure a valid SMTP gateway is specified (**Maintenance > Event Logging > Mail**). See [Section C: "Configuring Event Logging and Notification"](#) on page 1309 for more information.

Use the following procedure to modify the current settings.

To change the threshold and notification properties:

1. Select the **Maintenance > Health Monitoring** tab.
2. Select the tab for the metric you wish to modify.
 - To change the system resource metrics, select **General**.
 - To change the hardware, ADN status and health check status metrics, select **Status**.
 - To change the licensing metrics, select **Licensing**.
3. Click **Edit** to modify the threshold and notification settings. The **Edit Health Monitor Setting** dialog displays. Hardware, health check, and ADN thresholds cannot be modified.



4. Modify the threshold values:
 - a. To change the critical threshold, enter a new value in the Critical Threshold field.


- b. To change the critical interval, enter a new value in the Critical Interval field.
 - c. To change the warning threshold, enter a new value in the Warning Threshold field.
 - d. To change the warning interval, enter a new value in the Warning Interval field.
5. Modify the notification settings.
 - **Log** adds an entry to the Event log.
 - **Trap** sends an SNMP trap to all configured management stations.
 - **Email** sends an e-mail to the addresses listed in the Event log properties. See [Section C: "Configuring Event Logging and Notification"](#) on page 1309 for more information.
6. Click **OK** to close the Edit Metric dialog.
7. Click **Apply**.

Viewing Health Monitoring Statistics

While the **Health:** indicator presents a quick view of the appliance's health, the **Statistics > Health Monitoring** page provides more information about the current state of the health monitoring metrics.


To review the health monitoring statistics:

1. From the Management Console, select **Statistics > Health Monitoring**.



2. Select a health monitoring statistics tab:
 - **General:** Lists the current state of CPU utilization, interface utilization, memory utilization, and cloud common policy errors.
 - **Licensing:** Lists the current state of license utilization and license expiration.

- **Status:** Lists the current state of ADN status, hardware (including disk status, temperature, fan speed, power supply) and health check status.
3. To get more details about a metric, highlight the metric and click **View**. The **View Metrics Detail** dialog displays.



4. Click **Close** to close the **View Metrics Detail** dialog.
5. Optional—To modify a metric, highlight the metric and click **Set Thresholds**. The **Maintenance > Health Monitoring** page displays. To modify the metric, follow the procedure describe in "[Changing Threshold and Notification Properties](#)" on page 1347.

The show system-resource-metrics command lists the state of the current system resource metrics.

See Also:

- ❑ "[About the General Metrics](#)" on page 1340
- ❑ "[About the Licensing Metrics](#)" on page 1341
- ❑ "[About the Status Metrics](#)" on page 1343

Interpreting Health Monitoring Alerts

If you need assistance with interpreting the health monitoring alerts you receive, contact Symantec Technical Support. For non-technical questions such as licensing or entitlements, contact Symantec Customer Support.

Symantec recommends the following guidelines to meet your support needs:

1. Consult articles and documentation at MySymantec:
https://support.symantec.com/en_US.html
2. (MySymantec login required) If your request is not urgent, open a support case at:
<https://mysymantec.force.com/customer/s/>

-
3. If your request is urgent, contact us:
<https://www.symantec.com/contact-us>

See Also:

- ❑ "About Health Monitoring" on page 1335
- ❑ "Planning Considerations for Using Health Monitoring" on page 1338
- ❑ "About the Health Monitoring Metric Types" on page 1338
- ❑ "About the General Metrics" on page 1340
- ❑ "About the Licensing Metrics" on page 1341
- ❑ "About the Status Metrics" on page 1343
- ❑ "Snapshot of the Default Threshold Values and States" on page 1345

Chapter 73: Verifying the Health of Services Configured on the ProxySG

This section discusses Blue Coat health checks, which enable you to determine the availability of external networking devices and off-box services.

Topics

Refer to the following topics:

- ❑ [Section A: "Overview" on page 1354](#)
- ❑ [Section B: "About Blue Coat Health Check Components" on page 1357](#)
- ❑ [Section C: "Configuring Global Defaults" on page 1363](#)
- ❑ [Section D: "Forwarding Host and SOCKS Gateways Health Checks" on page 1373](#)
- ❑ [Section E: "DNS Server Health Checks" on page 1377](#)
- ❑ [Section F: "Authentication Health Checks" on page 1380](#)
- ❑ [Section G: "Virus Scanning and Content Filtering Health Checks" on page 1382](#)
- ❑ [Section H: "Managing User-Defined Health Checks" on page 1385](#)
- ❑ [Section I: "Viewing Health Check Statistics" on page 1392](#)
- ❑ [Section J: "Using Policy" on page 1397](#)

Section A: Overview

The ProxySG performs health checks to test for network connectivity and to determine the responsiveness of external resources. Examples of external resources include: DNS servers, forwarding hosts, SOCKS gateways, authentication servers, and ICAP services (for example, anti-virus scanning services).

The ProxySG automatically generates health checks based on:

- ❑ Forwarding configuration
- ❑ SOCKS gateways configuration
- ❑ DNS server configuration
- ❑ ICAP service configuration
- ❑ Authentication realm configuration
- ❑ Whether Dynamic Real-Time Rating (WebPulse) is enabled

You also can create user-defined health checks, including a composite health check that combines the results of multiple other health check tests. For information on health check types, see [Section B: "About Blue Coat Health Check Components"](#) on page 1357.

Health checks fall into three broad categories:

- ❑ Determining if the IP address can be reached. Health check types that fall into this category are:
 - Forwarding hosts
 - SOCKS gateways
 - User-defined host health checks
- ❑ Determining if a service is responsive. Health check types that fall into this category are:
 - Authentication servers
 - DNS server
 - Dynamic Real-Time Rating (WebPulse) service
 - ICAP services

- ❑ Determining if a group is healthy. Group tests are compilations of individual health checks, and the health of the group is determined by the status of the group members. Health check types that fall into this category are:
 - Forwarding groups
 - SOCKS gateway groups
 - ICAP service groups
 - User-defined composite health checks

Information provided by health checks allows you to accomplish the following:

- ❑ Detect potential network issues before they become critical. For example, if the health check for an individual host fails, the ProxySG sends an alert (using e-mail, SNMP, or by writing to an event log) to the designated recipients, if configured. To configure recipients, see [Section C: "Configuring Global Defaults"](#) on page 1369.
- ❑ Track response times and report failures. For example, if the DNS server performance suffers a reduction, the users experience response time delays. The DNS health check records the average response time (in milliseconds) and allows you to interpret the reason for the performance reduction. Should the DNS server become unavailable, the failed health check triggers an alert.

Furthermore, the ProxySG uses health check information to accomplish the following:

- ❑ When combined with failover configurations, health checks redirect traffic when a device or service failure occurs. For example, a health check detects an unhealthy server and a forwarding rule redirects traffic to a healthy server or proxy.
- ❑ Monitor the impact of health check states on the overall health of the appliance. Health check status is a metric in calculating the overall health of the ProxySG and is reflected in the health monitor, which is located at the upper right hand corner of the Management Console. For example, if a health check fails, the health monitor displays **Health: Warning**. You can click on the health monitor link to navigate and view the cause for the warning.

Executing an instant health check

Although the ProxySG automatically executes health checks, you can perform an instant health check from the **Configuration > Health Checks > General > Health Checks** tab by selecting the health check and clicking **Perform health check**. You can also view the health check state on the **Statistics > Health Check** tab.

Background DNS Resolution


Background testing of the DNS resolutions is performed on all resolvable hostnames used in the health check system, including forwarding hosts, WebPulse, and SOCKS gateways. That way, the list of IP addresses associated with a hostname stays current. The DNS system is checked whenever the time-to-live (TTL) value of the DNS entry expires.

Note: If a hostname consists of a dotted IP address, no DNS resolution is performed.

When a host is resolved by DNS to multiple IP addresses, health checks keep those addresses current through background updates. You can specify the timing for the updates. After the test or tests are conducted for each IP address, the results are combined. If the result for any of the resolved IP addresses is healthy, then the host is considered healthy because a healthy connection to that target can be made.

To specify the intervals for background DNS testing:

1. Select the **Configuration > Health Checks > Background DNS** tab.

The screenshot shows a configuration window titled "Background DNS". Inside, there is a sub-section titled "Background DNS Update". It contains three settings: "Minimum time to live for DNS results:" with a value of 30 seconds; "Maximum time to live for DNS results:" with a checked checkbox and a value of 300 seconds; and "Interval to wait after DNS resolution error:" with a value of 10 seconds. Each value is in a text input field followed by the word "seconds".

2. Specify options, as necessary:
 - a. Minimum time to live for DNS results—Cannot be zero (0). Test results are valid for this length of time. Retests can occur any time after this value.
 - b. Maximum time to live for DNS results—(Optional) How long the DNS test results remain valid before a retest is required.
 - c. Interval to wait after DNS resolution error—If the background DNS test discovers errors, this value specifies how long to wait before retesting. If a specific error repeatedly displays in the event logs, further network troubleshooting is required.
3. Click **Apply**.

Section B: About Blue Coat Health Check Components

Health checks have two components:

- ❑ Health check type: The kind of device or service the specific health check tests. The following types are supported:
 - Forwarding host and forwarding group
 - SOCKS gateway and SOCKS gateway group
 - DNS servers
 - External Authentication servers
 - ICAP service and ICAP service group
 - Dynamic Real-Time Rating Service
 - User-defined host and composite health checks
- ❑ Health check tests: The method of determining network connectivity, target responsiveness, and basic functionality.
 - Health checks (external targets)
 - Authentication
 - Internet Control Message Protocol (ICMP)
 - DNS
 - TCP
 - SSL
 - HTTP
 - HTTPS
 - ICAP
 - WebPulse
 - Health checks (group targets)
 - Groups
 - Composite

Note: Some health checks (such as forwarding hosts and SOCKS gateways) can be configured to report the result of a composite health check instead of their own test.

Some health check types only have one matching test, while others have a selection. For more information about health check types and tests, see [Table 73–1](#) on page 1359.

About Health Check Types

Most health checks are automatically created and deleted when the underlying entity being checked is created or deleted. When a forwarding host is created, for example, a health check for that host is created. Later, if the forwarding host is deleted, the health check for it is deleted as well. User interaction is not required, except to change or customize the health check behavior if necessary. However, if a health-check is referenced in policy, you cannot delete the corresponding host or the health check itself until the reference in policy is deleted.

In addition to the automatically generated health checks generated, run, and deleted, Blue Coat also supports two types of user-defined health checks. These health checks are manually created, configured, and deleted.

- ❑ *Composite* health checks: A method to take the results from a set of health checks (automatically generated or user-defined health checks) and combine the results.
- ❑ *Host* health checks: A method to test a server, using a selection of ICMP, TCP, SSL, HTTP, and HTTPS tests.

Note: Although a host health check tests an upstream server, it can also be used to test whether a proxy is working correctly. To test HTTP/HTTPS proxy behavior, for example, you can set up a host beyond the proxy, and then use forwarding rules so the health check passes through the proxy to the host, allowing the proxy to be tested.

User-defined health checks allow you to test for attributes that the ProxySG does not test automatically. For example, for a forwarding host, you could perform three user-defined tests — an HTTP test, an HTTPS test, and a TCP test of other ports. Then, you can set up a composite health check that combines the results of these user-defined tests to represent the health of the forwarding host. The ProxySG reports the status of the (user-defined) composite health check as the forwarding host's health, instead of the default forwarding host health check.

All health check types are given standardized names, based on the name of the target. For example:

- ❑ Forwarding hosts and groups have a prefix of **fwd**
- ❑ DNS servers have a prefix of **dns**
- ❑ SOCKS gateways and gateway groups have a prefix of **socks**
- ❑ Authentication realms have a prefix of **auth**
- ❑ External services have prefixes of **icap**, and **WebPulse**
- ❑ User-defined or composite health checks have a prefix of **user**

Health Check Tests

Based on the health check type, the ProxySG periodically tests the health status, and thus the availability of the host. You can configure the time interval between tests. If the health check test is successful, the appliance considers the host available.

The health check tests are described in the table below.

Table 73–1 Health Check Tests

Health Check Test	Description	Used With Health Check Type
Response Times	The minimum, maximum, and average response times are tracked, with their values being cleared whenever the health check changes state.	All
ICMP Test (Layer 3)	The basic connection between the ProxySG and the origin server is confirmed. The server must recognize ICMP echoing, and any intervening networking equipment must support ICMP. The ProxySG appliance sends a ping (three ICMP echo requests) to the host. ICMP tests do not support policy for SOCKS gateways or forwarding.	Forwarding hosts, SOCKS gateways, or user-defined hosts
TCP Socket Connection Test (Layer 4)	A TCP test establishes that a TCP layer connection can be established to a port on the host. Then the connection is dropped. TCP tests for a SOCKS gateway do not support policy for SOCKS gateways or forwarding. TCP tests for a forwarding host or a user-defined health check support SOCKS gateways policy but not forwarding policy.	Forwarding hosts, SOCKS gateways, or user-defined hosts
SSL Test	A connection is made to a target and the full SSL handshake is conducted. Then, much like the TCP test, the connection is dropped. For a forwarding host, a terminating HTTPS port must be defined or the test fails. SSL tests for a forwarding host or a user-defined health check support SOCKS gateways policy. The SSL tests do not support forwarding policy. An SSL test executes the SSL layer in policy and obeys any settings that apply to server-side certificates, overriding any settings obtained from a forwarding host.	Forwarding hosts or user-defined hosts

Table 73–1 Health Check Tests (Continued)

Health Check Test	Description	Used With Health Check Type
HTTP/HTTPS Tests for Servers and Proxies	HTTP/HTTPS tests execute differently depending on whether the upstream target is a server or a proxy. For a forwarding host, the server or a proxy is defined as part of the forwarding host configuration. For a user-defined health check, the target is always assumed to be a server. For a server: • The HTTP test sends an HTTP GET request containing only the URL path to an HTTP port. • The HTTPS test sends an HTTPS GET request containing only the URL path over an SSL connection to a terminating HTTPS port. If an appropriate port is not available on the target, the test fails. For a proxy: • The HTTP test sends an HTTP GET request containing the full URL to an HTTP port. • Since a server is required to terminate HTTPS, the HTTPS test sends an HTTP CONNECT request to the HTTP port. If an appropriate HTTP port is not available on the proxy, either test fails. An HTTP/HTTPS test requires a full URL for configuration. The HTTP/HTTPS tests for a forwarding host support SOCKS gateway policy but not forwarding policy. The HTTP/HTTPS tests for a user-defined health check support SOCKS gateway and forwarding policy. An HTTPS test executes the SSL layer in policy and obeys any settings that apply to server-side certificates, overriding any settings obtained from a forwarding host.	Forwarding hosts or user-defined hosts.
HTTP/HTTPS Authentication	For HTTP/HTTPS tests, you can test authentication using a configured username and password. The passwords are stored securely in the registry.	Forwarding hosts or user-defined hosts.

Table 73–1 Health Check Tests (Continued)

Health Check Test	Description	Used With Health Check Type
HTTP/HTTPS Allowed Responses	For an HTTP or HTTPS test, this is the set of HTTP response codes that indicate success. The default is to accept only a 200 response as successful. You can specify the sets of response codes to be considered successful.	Forwarding hosts or user-defined hosts.
External ServicesTests	The tests for external services are specialized tests devised for each particular kind of external service. The health check system conducts by sending requests to the configured services, which reports back a health check result.	ICAP, off-box, WebPulse services.
Group	Individual tests that are combined for any of the four different available groups (forwarding, SOCKS gateways, and ICAP services). If any of the members is healthy, then the group as a whole is considered healthy. Note: Symantec supports a composite test, used only with composite (user-defined) health checks, that is similar to a group test except that, by default, all members must be healthy for the result to be healthy. These settings are configurable. By default, group health tests are used for two purposes: • Monitoring and notification • Policy	Forwarding groups, SOCKS gateways groups, and ICAP external service groups.
DNS Server	The DNS server maps the hostname, default is www.bluecoat.com, to an IP address. The health check is successful if the hostname can be resolved to an IP address by the DNS server.	DNS
Authentication	Authentication health checks assess the realm's health using data maintained by the realm during active use. Authentication health checks do not probe the authentication server with an authentication request.	Authentication

See Also

- ❑ ["To edit forwarding and SOCKS gateways health checks:"](#) on page 1374
- ❑ ["To edit forwarding or SOCKS gateway group health checks:"](#) on page 1375
- ❑ ["To edit a DNS server health check:"](#) on page 1378
- ❑ ["To edit an authentication health check:"](#) on page 1380
- ❑ ["To edit virus scanning and content filtering tests:"](#) on page 1382
- ❑ ["To edit ICAP group tests:"](#) on page 1383
- ❑ ["To create a user-defined host health check:"](#) on page 1387
- ❑ ["Click Apply.To create a user-defined composite health check:"](#) on page 1389

Section C: Configuring Global Defaults

All health checks are initially configured to use global defaults. The only exception is the *Dynamic Categorization* service, which has the healthy interval set to 10800 seconds (3 hours), and the failure trigger set to 1.

About Health Check Defaults

You can change the defaults on most health checks. These defaults override global defaults, which are set from the **Configuration > Health Checks > General > Default Settings** tab.

You can edit health check intervals, severity, thresholds, and notifications for automatically generated health checks in two ways:

- ❑ Setting the global defaults. These settings affect all health checks, unless overridden by explicit settings.
- ❑ Setting explicit values on each health check.

The default health check values are:

- ❑ Ten seconds for healthy and sick *intervals* (an interval is the period between the completion of one health check and the start of the next health check).
- ❑ One for healthy and sick *thresholds*. A healthy threshold is the number of successful health checks before an entry is considered healthy; a sick threshold is the number of unsuccessful health checks before an entry is considered sick.
- ❑ Warning for the *severity* notification, which governs the effect that a health check has on the overall health status of the ProxySG.
- ❑ Disabled for logging health check status using e-mails, event logs, or SNMP traps.

To configure the settings, continue with "[Changing Health Check Default Settings](#)" on page 1366.

To configure notifications, continue with "[Configuring Health Check Notifications](#)" on page 1369.

Enabling and Disabling Health Checks

You can enable or disable health checks and configure them to report as healthy or unhealthy during the time they are disabled.

Setting a health check as disabled but reporting healthy allows the ProxySG to use the device or service without performing health checks on it. If, for example, you have configured a forwarding host on the ProxySG, a health check for the forwarding host is automatically created. If you then configure the health check as disabled reporting healthy, the ProxySG considers the forwarding host as healthy without performing periodic health checks on it.

If the case of a group health check that is disabled but reporting healthy, all members of the group are treated as healthy regardless of the status of the members' individual health check result.

Note: Individual health checks for members of a group remain active; they can be used apart from the group.

Setting a health check as disabled but reporting sick is useful to remove an upstream device for servicing, testing, or replacement. This setting takes the device offline after it completes processing pre-existing traffic. Then the device can be safely disconnected from the network without altering any other configuration.

You cannot enable or disable all health checks at once.

To enable a health check:

1. In the Management Console, select **Configuration > Health Checks > General**.
2. On the Health Checks tab, select the health check you want to enable.
3. Click **Edit**.
4. For Enabled state, select **Enabled** and then click **OK**.
5. Click **Apply**.

To disable a health check:

1. In the Management Console, select **Configuration > Health Checks > General**.
2. On the Health Checks tab, select the health check you want to disable.
3. Click **Edit**.
4. For Enabled state, select one of the following:
 - To report the health check as healthy, select **Disabled: Healthy**.
 - To report the health check as unhealthy, select **Disabled: Unhealthy**.
5. Click **OK**.
6. Click **Apply**.

Notifications and SNMP Traps

If you configure notifications, the ProxySG sends all or any of e-mail, SNMP, and event log notifications when a change of health check state occurs. By default, all notifications are disabled.

On the ProxySG you can:

- ☐ Globally change notifications for all health checks
- ☐ Explicitly change notifications for specific health checks
- ☐ Enable notifications of transitions to healthy
- ☐ Enable notifications of transitions to unhealthy

A transition to healthy occurs as soon as the target is sufficiently healthy to be sent a request, even though the target might not be completely healthy. For example, if you have multiple IP addresses resolved and only one (or a few) is responsive, the group is classified as healthy and the health status might be **Ok with errors** or **Ok for some IPs**. For some health check groups, like forwarding hosts, you can configure a minimum number of members that must be healthy for the group to be healthy.

In the event log, status changes can be logged as either informational or severe logs. In addition to the overall health of the device, you can enable notifications for each resolved IP address of a target device (if applicable).

An SNMP trap can also be used for notification of health check state changes. It is part of the Blue Coat Management Information Base (MIB) as *blueCoatMgmt 7.2.1*. For information on configuring SNMP, see "[Monitoring Network Devices \(SNMP\)](#)" on page 1319.

Guidelines for Setting the Severity of a Health Check

Severity indicates how a failed health check affects the overall health of the device. The **severity** option links **Health Checks** and **Health Monitoring**. The health monitor displays the overall health of the device after considering the health check status in conjunction with other health monitoring metrics. For information on the health monitoring metrics, see "[Configuring Health Monitoring](#)" on page 1335.

Note: Severity of a health check is pertinent only when a health check fails.

The ProxySG allows you to configure the severity option to Critical, Warning and No effect. Set the severity of a health check to:

- ❑ **Critical:** If the success of a health check is crucial to the health of the device. If the health check then reports unhealthy, the overall health status becomes **Critical**.
- ❑ **Warning:** If a failed health check implies an emerging issue and the administrator must be alerted when the health check state transitions from healthy to unhealthy. Consequently, when the health check reports unhealthy, the overall health status transitions to **Warning**.
- ❑ **No effect:** If the success of a health check bears no impact on the health of the device. Should the health check transition to unhealthy, the overall health status of the device retains its current status and does not change.

For example, if the severity on an external service health check for **ICAP**, is set to severity level **Critical** and the health check fails, the overall health status of the device will transition to **Health: Critical**.

To change notifications, continue with "[Configuring Health Check Notifications](#)" on page 1369.


Notification E-mail Contents

When the health status of the ProxySG changes (based on the health check parameters) a notification e-mail is sent to the user(s) on the event logging notification list. The notification e-mail contains information relevant to the health check test that has been triggered. The information can be used as reference information or to troubleshoot a variety of errors.

Note: E-mail notifications are turned off by default. To enable e-mail notifications, see "[Configuring Health Check Notifications](#)" on page 1369.

When a status change notification e-mail is sent to a listed user, it includes the following information in the e-mail subject line:

- ❑ Appliance name (see the *Initial Configuration Guide* for more information on naming an appliance)
- ❑ Health check test (see "[Health Check Tests](#)" on page 1359 for a list of available tests)
- ❑ Health state change (Health state changes are contingent upon health check parameters)



Health check 'Proxy_name' changed from: Check failed, to: OK, and has become useable. Status: Success. (IP 10.169.0.19)

The body of the e-mail includes relevant information based on the nature of the health change.

Changing Health Check Default Settings

You can modify the default settings for all health checks on the **Configuration > Health Checks > General > Default Settings** tab or you can override the default settings for a health check on the **Configuration > Health Checks > General > Health Checks** tab, selecting the health check, and clicking **Edit**. Explicit health settings override the global defaults.

To change the global default settings:

1. Select **Configuration > Health Checks > General > Default Settings**.

Health Checks	Default Settings	Default Notifications
Default Settings		
Healthy interval:	10	seconds
Healthy threshold:	1	
Sick interval:	10	seconds
Sick threshold:	1	
☐ Failure trigger threshold:	0	
☐ Response time threshold:	0	milliseconds


2. Change the settings as appropriate:
 - a. Specify the healthy interval, in seconds, between health checks. The default is **10**. The healthy interval can be between 1 second and 31536000 seconds (about one year).
 - b. Specify the healthy threshold for the number of successful health checks before an entry is considered healthy. Valid values can be between 1 and 65535. The default is **1**.
 - c. Specify the sick interval, in seconds, between health checks to the server that has been determined to be unhealthy or out of service. The default is **10**. The sick interval can be between 1 second and 31536000 seconds (about 1 year).
 - d. Specify the sick threshold, or the number of failed health checks before an entry is considered unhealthy. Valid values can be between 1 and 65535. The default is **1**.
 - e. Specify the failure threshold for the number of failed connections to the server before a health check is triggered. Valid values can be between 1 and 2147483647. It is disabled by default.

The failures are reported back to the health check as a result of either a connection failure or a response error. The number of these external failures is cleared every time a health check is completed. If the number of failures listed meets or exceeds the threshold and the health check is idle and not actually executing, then the health of the device or service is immediately checked.

- f. Specify the maximum response time threshold, in milliseconds. The threshold time can be between 1 and 65535.
3. Click **Apply**.

To override default settings for a targeted health check:

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Select the test you want to modify.
3. Click **Edit**. The example below uses a SOCKS gateway.



4. To substitute special values for this test:
 - a. Click **Override the default settings**. The Override Default Settings dialog displays. Configure the override options. You can cancel your choices by clicking **Clear all overrides**.
 - b. Specify the healthy interval, in seconds, between health checks to the server. The default is **10**. The healthy interval is between 1 second and 31536000 seconds (about one year).
 - c. Specify the healthy threshold for the number of successful health checks before an entry is considered healthy. Valid values are 1-65535. The default is **1**.
 - d. Specify the sick interval, in seconds, between health checks to the server that has been determined to be unhealthy or out of service. The default is **10**. The sick interval is between 1 second and 31536000 seconds (about 1 year).

- e. Specify the sick threshold, or the number of failed health checks before an entry is considered unhealthy. Valid values are 1-65535. The default is **1**.
- f. Specify the failure trigger for the number of failed connections to the server before a health check is triggered. Valid values are between 1 and 2147483647.

The failures are reported back to the health check as a result of either a connection failure or a response error. The number of these external failures is cleared every time a health check is completed. If the number of failures listed meets or exceeds the threshold, and the health check is idle and not actually executing, then the health of the device or service is immediately checked.

- g. Specify the maximum response time threshold, in milliseconds. The threshold time can be between 1 and 65535.
 - h. Click **OK** to close the dialog.
5. Click **Apply**.

Configuring Health Check Notifications

The ProxySG allows you to configure notifications that alert you to changes in health status and to emerging issues. By default, notifications for health check events and status are disabled.

You can set up health check notifications:

- ❑ Globally on the **Configuration > Health Checks > General > Default Notifications** tab
- ❑ Explicitly, for a health check, on the **Configuration > Health Checks > General > Health Checks** tab, selecting the health check, and clicking **Edit**.
Explicit health settings override the global defaults.

To configure health check notifications globally:


1. Select **Configuration > Health Checks > General > Default Notifications**.
2. Select the **Severity** level for the health check.
 - Critical: If the health check fails, the device is in critical condition
 - Warning: If the health check fails, the device needs to be monitored and the health check status displays as Warning. This is the default setting.
 - No effect: The health check has no impact on the overall health of the device.

The screenshot shows the 'Default Notifications' configuration window. It has three tabs: 'Health Checks', 'Default Settings', and 'Default Notifications'. The 'Default Notifications' tab is active. Inside, there's a section titled 'Default Notifications' with a blue header. Below it, the 'Severity' is set to 'No effect' via a dropdown. The 'E-mail notification' section has three checkboxes: 'Notify on a transition to healthy' (unchecked), 'Notify on a transition to sick' (unchecked), and 'Notify for each IP address' (checked). The 'Event logging' section has two dropdowns for logging transitions: 'Log a transition to healthy as:' (set to 'Informational') and 'Log a transition to sick as:' (set to 'Severe'), plus a checked checkbox for 'Log for each IP address'. The 'SNMP traps' section has three unchecked checkboxes for different transition states.

3. Select the options to enable notifications:
 - a. **E-mail notification:** Select the appropriate check boxes to enable the e-mail notifications you require. Recipients are specified in **Maintenance > Event Logging > Mail**.
 - b. **Event logging:** Select the appropriate options to enable the event logging you require. Messages can be logged as either informational or severe.
 - c. **SNMP traps:** Select the situations for which you require SNMP traps to be sent.
4. Click **Apply**.

To override the default notifications for a targeted health check:

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Select a test to modify.
3. Click **Edit**. The Edit dialog displays. The following example uses a forwarding host.
4. To change default notifications for this test, select **Override the default notifications**. By default, notifications are not sent for any health checks.



5. Select the options to override. You can cancel your choices by clicking **Clear all overrides**.
 - a. **Severity:** Select the severity option as required.
 - **Critical:** If the health check fails, the device is in critical condition
 - **Warning:** If the health check fails, the device needs to be monitored and the health check status displays as Warning. This is the default setting.
 - **No effect:** The health check has no impact on the overall health of the device.
 - b. **Override E-mail notification:** Select the appropriate check boxes to enable the e-mail notifications you require. Specify recipients in **Maintenance > Event Logging > Mail**.

- c. **Event logging:** Select the appropriate check boxes to enable the event logging you need. Messages can be logged as either informational or severe.
 - d. **SNMP traps:** Select the situations in which you want SNMP traps to be sent.
 - e. Click **OK** to close the override dialog
 - f. Click **OK** to close the edit dialog.
6. Click **Apply**.

Section D: Forwarding Host and SOCKS Gateways Health Checks

Before you can edit forwarding or SOCKS gateways health check types, you must configure forwarding hosts or SOCKS gateways. For information about configuring forwarding, see [Chapter 43: "Configuring the Upstream Network Environment"](#) on page 865; for information about configuring SOCKS gateways, see [Chapter 41: "SOCKS Gateway Configuration"](#) on page 839.

This section discusses managing the automatically generated forwarding host and SOCKS gateway health checks.

About Forwarding Hosts and SOCKS Gateways Configurations

The forwarding host health check configuration defines whether the target being tested is a server or a proxy, which ports are available, and provides the setting for the server certificate verification.

The SOCKS gateways health check configuration defines the SOCKS port, the version (4 or 5), and possibly a username and password.

Forwarding Hosts Health Checks

The default for a newly created forwarding host is a TCP health check using the first port defined in the forwarding host's port array (typically the HTTP port). You can change the port setting. The TCP test can support SOCKS gateway policy. The URL uses the forwarding host hostname, such as:

```
tcp://gateway_name:port/
```

SOCKS Gateways Health Checks

The default for a newly created SOCKS gateway is a TCP health check using the SOCKS port in the SOCKS gateways configuration.

Forwarding and SOCKS Gateways Groups Health Checks

Specific tests are not done for groups. Health check test results are determined from examining and combining the health of the group members.

Note: You can create groups in the **Configuration > Forwarding > Forwarding Hosts** tab or **Configuration > Forwarding > SOCKS Gateways** tab.

By default, if any of the members of the group are healthy, then the group is considered healthy. You can specify the number of group members that must be healthy for the group to be considered healthy.


Editing Forwarding and SOCKS Gateways Health Checks

You can edit, but not delete, the forwarding and SOCKS gateway tests and groups. The settings you can change are:

- ❑ Enable or disable the health check
- ❑ Override default notifications
- ❑ Select the type of test
- ❑ Specify settings for the selected test
- ❑ Override default settings
- ❑ Select the minimum number of healthy members for a group to report healthy

To edit forwarding and SOCKS gateways health checks:

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Select the forwarding host test or SOCKS gateways test to modify.
3. Click **Edit**.



4. Make the necessary changes:
 - a. Select the **Type of Test** from the drop-down list.
 - b. Select the **Enabled state** radio button as required.

- c. Select the port setting you require. If you select **Use Port**, enter the new port number.
 - d. To change the default settings for this test, click **Override the default settings**. Select the options to override. Cancel your choices by clicking **Clear all overrides**. For detailed information about configuring healthy and sick intervals and thresholds, see ["Changing Health Check Default Settings"](#) on page 1366. Click **OK** to close the dialog.
 - e. To change default notifications, click **Override the default notifications**. By default, no notifications are sent for any health checks. Select the options to override. You can cancel your choices by clicking **Clear all overrides**. For detailed information about configuring notifications, see ["Configuring Health Check Notifications"](#) on page 1369. Click **OK** to close the dialog.
 - f. Click **OK** to close the edit dialog.
5. Click **Apply**.

To edit forwarding or SOCKS gateway group health checks:

Note: The only way to add or delete group members to the automatically generated health check tests is to add and remove members from the actual forwarding or SOCKS gateway group. The automatically generated health check is then updated.

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Select the forwarding or SOCKS gateways group health check you need to modify.
3. Click **Edit**.

The screenshot shows the 'Edit Health Check Group' dialog box. The 'Name' field is set to 'fwd.test2'. The 'Enabled state' is set to 'Enabled'. The 'Minimum number of members that must be healthy for group to be healthy' is set to '1'. There is an 'Override Defaults' section with an 'Override' button and the text 'the default notifications'. The 'OK' and 'Cancel' buttons are at the bottom.

4. Make the necessary changes:
 - a. Select an **Enabled state** option.
 - b. Select the **Minimum number of users that must be healthy for group to be healthy** from the drop-down list.
 - c. To create notification settings, click **Override the default notifications**. Select the options. Cancel your choices by clicking **Clear all overrides**. For detailed information about configuring notifications, see ["Configuring Health Check Notifications"](#) on page 1369.
 - d. Click **OK** to close the override dialog.
 - e. Click **OK** to close the health check group.
5. Click **Apply**.

Section E: DNS Server Health Checks

- ❑ ["About DNS Server Health Checks"](#)
- ❑ ["Editing DNS Server Health Checks"](#)

About DNS Server Health Checks

A DNS server health check is automatically generated for each DNS server configured on the ProxySG and is deleted when the DNS server is removed. For information on configuring DNS servers, see ["Adding DNS Servers to the Primary or Alternate Group"](#) on page 816.

The ProxySG uses DNS server health checks to verify the responsiveness of the DNS server. The health check status is recorded as:

- ❑ **Healthy**, when the ProxySG successfully establishes a connection with the DNS server and is able to resolve the configured hostname.
- ❑ **Unhealthy**, either if the ProxySG is unable to establish a connection with the DNS server, or if the ProxySG is unable to resolve the configured hostname. The status reports **Check failed** or **DNS failed**.

When a DNS server is unhealthy, the ProxySG avoids contacting that server and directs requests to other DNS servers configured in the group, as applicable.

The DNS health check attempts to look up a configurable hostname. The default hostname depends on the DNS configuration:

- ❑ For a server in the primary or alternate DNS group, the default is `www.bluecoat.com`.
- ❑ For a server in a custom DNS group, the default is the longest domain name listed in the group.

You can also override these defaults and specify a health check hostname for each DNS server.

See Also

[Chapter 37: "Configuring DNS"](#) on page 813


Editing DNS Server Health Checks

On the ProxySG, you can edit the following settings for a DNS server health check:

- ❑ Enable or disable the health check
- ❑ Specify a hostname
- ❑ Override default settings — change healthy and sick intervals, and thresholds
- ❑ Override default notifications — change the severity and notification options for alerts

To edit a DNS server health check:

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Select the DNS health check to modify.
3. Click **Edit**. The Edit DNS server dialog displays.



4. Configure the DNS server health check options:
 - a. Select the **Enabled state** option, as required.
 - **Enabled** allows the ProxySG to query the DNS server and to report changes in the health state.
 - **Disabled, reporting as healthy** disables the health check and reports the service as healthy.
 - **Disabled, reporting as sick** disables the health check and reports the service as unhealthy.

- b. Select the **Host** option, as required.
 - **Use default host** uses the default hostname.
 - **Use user defined host** allows you to configure a custom hostname for this health check. Enter the hostname in the box provided.

Proceed to Step e if you do not want to override defaults.

- c. To change default settings, click **Override the default settings**. Select the options to override. Cancel your choices by clicking **Clear all overrides**. For detailed information about configuring healthy and sick intervals and thresholds, see "[Changing Health Check Default Settings](#)" on page 1366. Click **OK** to close the dialog.
 - d. To change the default notifications, click **Override the default notifications**. Select the options. Cancel your choices by clicking **Clear all overrides**. For detailed information about configuring notifications, see "[Configuring Health Check Notifications](#)" on page 1369.
 - e. Click **OK** to close the override dialog.
5. Click **OK** to close the edit dialog.
 6. Click **Apply**.

Section F: Authentication Health Checks

This section includes information on authentication server health checks. For information on authentication realms, see ["Controlling User Access with Identity-based Access Controls"](#) on page 898.

An authentication health check is automatically generated for each external authentication realm that is configured on the ProxySG. Authentication health checks assess the realm's health based on data gathered during the most recent authentication attempt. The response time recorded for this health check represents the average response time between two consecutive health checks.

Unlike most health checks, authentication health checks do not probe the target realm with an authentication request. Therefore, the health check will report healthy until the ProxySG records a failed authentication attempt.

The health states for authentication health checks can be:

- ❑ **Ok**, when the ProxySG records successful authentication attempts.
- ❑ **Check failed**, when the device records an unsuccessful authentication attempt.
- ❑ **Functioning on alternate server**, when a realm is operating on its alternate server.
- ❑ **Functioning properly with errors**, when the health check records intermittent failures on a server.

On an authentication health check, you can edit the following settings:


- ❑ Enable or disable the health check
- ❑ Override default settings — change healthy and sick intervals, and thresholds
- ❑ Override default notifications — change the severity, and notification options for alerts

By default, the health check is enabled and the ProxySG tracks the response time for the most recent authentication attempts. The other options are — Disabled, reporting sick and Disabled, reporting healthy.

Use the Disabled, reporting sick option when an authentication server requires downtime for maintenance, or the server is taken off-line temporarily. And the Disabled, reporting healthy option is relevant when you elect to use an authentication server despite failures in authentication attempts.

To edit an authentication health check:

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Select the `auth.test_name` health check to modify.
3. Click **Edit**. The Edit Authentication health check dialog displays.



4. Configure the authentication health check options:
 - a. Select the **Enabled state** radio button as required.
 - b. To change the default settings, click **Override the default settings**. Select the options to override. Cancel your choices by clicking **Clear all overrides**. For detailed information about configuring healthy and sick intervals and thresholds, see ["Changing Health Check Default Settings"](#) on page 1366. Click **OK** to close the dialog.
 - c. To change the default notifications, click **Override the default notifications**. Select the options. Cancel your choices by clicking **Clear all overrides**. For detailed information about configuring notifications, see ["Configuring Health Check Notifications"](#) on page 1369. Click **OK** to close the dialog.
 - d. Click **OK** to close edit dialog.
5. Click **Apply**.

Section G: Virus Scanning and Content Filtering Health Checks

The virus scanning and content filtering services include ICAP services and WebPulse. While these health checks are created and deleted automatically, the service itself must be created before health checks can be used. For more information about creating ICAP services, see [Chapter 20: "Filtering Web Content"](#) on page 377. The WebPulse service health check is automatically created if you use Blue Coat WebFilter and the rating service is enabled.

The health check system conducts external service tests by sending requests to each configured service and reports back a health check result. The tests for each service is specialized and is devised specifically for each type of service.

Note: The names of the ICAP services and service groups can be a maximum of 64 characters long, a change from previous releases, which allowed names to be a maximum of 127 characters.

The settings you can change on ICAP, and WebPulse service health checks are:

- ☐ Enable or disable the health check
- ☐ Override default settings
- ☐ Override default notifications

To edit virus scanning and content filtering tests:

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Select the external service to modify. External services have prefix names of **WebPulse**, and **icap**.
3. Click **Edit**.
4. Make the necessary changes:
 - a. Select the **Enabled state** radio button as required.
 - b. To change default settings, click **Override the default settings**.
 - Select the check boxes to override. Cancel your choices by clicking **Clear all overrides**. For detailed information about configuring healthy and sick intervals and thresholds, see ["Changing Health Check Default Settings"](#) on page 1366.

Note: The WebPulse health check has default settings that differ from the defaults for other external services: 10800 seconds (3 hours) for the interval, and 1 for the failure trigger.

- Click **OK**.

- c. To change default notifications, click **Override the default notifications**. By default, no notifications are sent for any health checks. Select the options. Cancel your choices by clicking **Clear all overrides**. For detailed information about configuring notifications, see "[Configuring Health Check Notifications](#)" on page 1369.
 - d. Click **OK** to close the override dialog.
 - e. Click **OK** to close the edit dialog.
5. Click **Apply**.

To edit ICAP group tests:

Note: The only way to add or delete group members to the automatically generated health check tests is to add and remove members from the ICAP services. The automatically generated health check type is then updated.

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Select the external service group health check to modify. Groups are identified in the **Type** column.
3. Click **Edit**.

4. Make the necessary changes:
 - a. Enable or disable the **Enabled state** radio button as required.
 - b. Select the **Minimum number of members that must be healthy for group to be healthy** from the drop-down list. The default is set to one.
 - c. To create notification settings, click **Override the default notifications**. Select the options. Cancel your choices by clicking **Clear all overrides**. For detailed information about configuring notifications, see "[Configuring Health Check Notifications](#)" on page 1369.

- d. Click **OK** to close the override dialog.
 - e. Click **OK** to close the edit dialog.
5. Click **Apply**.

Section H: Managing User-Defined Health Checks

You can manually create and manage ICMP, TCP, HTTP, HTTPS, or SSL health check tests for any upstream TCP/IP device. You can use these user-defined health check types to send notifications of health check state changes.

Under most circumstances, you do not need to create user-defined health checks because the automatically generated health checks meet most needs. However, to check for things that Blue Coat does not test for automatically — for example, the health of the Internet or of the router, you might create user-defined health checks.

If, for example, you want to control Web traffic based on the apparent health of the Internet, you can create a user-defined health check to target known Internet sites. As long as a certain number of the sites are healthy, you can consider the Internet as healthy.

Further, you can use policy to configure forwarding rules on the ProxySG. Subsequently, if the user-defined health check determining internet accessibility transitions to unhealthy, all requests directed to the ProxySG will be forwarded to the alternate ProxySG until the primary ProxySG transitions to healthy again.

Note: Frequent testing of specific Internet sites can result in that Internet site objecting to the number of hits.

Symantec supports two types of user-defined health checks:

- ❑ **Host:** This health check type is for any upstream TCP/IP device. For more information, continue with "[About User-Defined Host Health Checks](#)".
- ❑ **Composite:** This health check type combines the results of other existing health checks. It can include other composite health checks, health checks for user defined hosts, and any automatically generated health checks. For more information, continue with "[About User-Defined Composite Health Checks](#)" on page 1386.

For information about configuring parameter and notification settings for automatically generated health check types, see [Section C: "Configuring Global Defaults"](#) on page 1363.

About User-Defined Host Health Checks

You can create, configure, and delete user-defined host health checks. These health checks support everything an automatically generated health check contains, including background DNS resolution monitoring and support for multiple addresses.

User-defined health checks can include:

- ❑ **ICMP:** The basic connection between the ProxySG and the origin server is confirmed. The server must recognize ICMP echoing, and any intervening networking equipment must support ICMP.

- ❑ TCP: Establishes that a TCP layer connection can be made to a port on the host. Then the connection is dropped.
- ❑ SSL: A connection is made to a target and the full SSL handshake is confirmed. Then the connection is dropped.
- ❑ HTTP/HTTPS: An HTTP or HTTPS test is defined by the URL supplied. The port used for this test is as specified in that URL. If no port is explicitly specified in the URL, the port defaults to the standard Internet value of 80 or 443.

When configuring user-defined host health check types, keep the following in mind:

- ❑ User-defined host health checks are created and deleted manually.
- ❑ All individual user-defined tests consider the target to be a server.
- ❑ To conduct proxy HTTP/HTTPS tests, a proxy must be defined as a forwarding host, set up between the originating device and the target, and forwarding policy must cause the test to be directed through the proxy.
- ❑ For an ICMP test, a hostname is specified in the health check configuration.
- ❑ The TCP and SSL tests support SOCKS gateway policy, based on a URL of `tcp://hostname:port/` and `ssl://hostname:port/`, respectively, using a hostname and port supplied in health check configuration.
- ❑ An HTTP/HTTPS test requires a full URL. The port used for this test is as specified in that URL. If no port is explicitly specified in the URL, the port defaults to the standard value for these protocols of 80 or 443. The server being tested is assumed to support whatever port is indicated.

Forwarding and SOCKS gateway policy is applied based on the URL. The HTTPS or SSL tests use all the server certificate settings in the SSL layer in policy. For a forwarding host, all the sever certificate settings in the SSL layer also apply, and if present, override the forwarding host configuration setting.

Note: None of the above tests apply to user-defined composite health checks, which only consist of a set of members and a setting to combine the results.

About User-Defined Composite Health Checks

You can create a composite health check to combine the results of multiple health checks. A composite health check can contain any number of individual health checks. Further, forwarding host and SOCKS gateway health checks can be configured to use the result of a composite health check.

By default, to report healthy, all members of a composite health check must be healthy. However, you can configure the number of members that must be healthy for the composite result to report healthy.

Composite health checks with no members always appear unhealthy.

Note: Automatically generated group tests and user-defined composite tests are not the same.

Group tests are automatically generated; they cannot be deleted. Some editing is permitted, but you cannot add or remove members of the group through the health checks module. You must modify the forwarding or SOCKS gateways groups to update the automatically generated group tests.

For a group test, the default is for the group to be healthy if any member is healthy. For a composite test, the default is for the group to be healthy if all members are healthy. (The default is configurable.)

Creating User-Defined Host and Composite Health Checks

You can create user-defined host and composite health checks for arbitrary targets.

Note: You cannot create user-defined health checks for external service tests, such as authentication servers, ICAP, and the WebPulse service.

The following procedure explains how to create a user-defined host health check. To create a user-defined composite health check, continue with "[Click Apply.To create a user-defined composite health check:](#)" on page 1389.

To create a user-defined host health check:

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Click **New**.

3. Select the type of test to configure from the **Type of test** drop-down list. To configure a composite test, see "[Click Apply.To create a user-defined composite health check:](#)" on page 1389.

The options you can select vary with the type of health check. The example above uses the HTTP/HTTPS options. Options for other tests are explained in this procedure, as well.

- a. Enter a name for the health check.
- b. Select the **Enabled state** option, as required.
- c. If you are configuring an SSL or TCP health check, enter the port to use.
- d. If you are configuring an ICMP, SSL, or TCP health check, enter the hostname of the health check's target. The hostname can be an IPv4 or IPv6 host or address.

- e. For HTTP/HTTPS only:
 - Enter the URL address of the target.
 - To use Basic user authentication, select the check box and enter the username and password of the target.
 - To use Basic proxy authentication because intermediate proxies might be between you and the target, select the check box and enter the username and password of the target.
 - To manage a list of HTTP/HTTPS response codes that are considered successes, enter the list in the **Allowed Response Code** field, separated by semi-colons. If one of them is received by the health check then the health check considers the HTTP(S) test to have been successful.

Note: The 200 response code is added by default. The list must always have at least one member.

- f. To change the default settings for this test, click **Override the default settings**. Select the override options. Cancel your choices by clicking **Clear all overrides**. For detailed information about configuring healthy and sick intervals and thresholds, see "[Changing Health Check Default Settings](#)" on page 1366. Click **OK**.
- g. To change the default notifications for this test, click **Override the default notifications**. By default, no notifications are sent for any health checks. Select the override options. You can cancel your choices by clicking **Clear all overrides**. For detailed information about configuring notifications, see "[Configuring Health Check Notifications](#)" on page 1369. Click **OK**.
- h. Click **OK** to close the dialog.

Click **Apply**. **To create a user-defined composite health check:**

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Click **New**.

3. Configure the options:
 - a. Select **Composite** from the **Type of Test** from the drop-down list.
 - b. Enable or disable the **Enabled state** option as required.
 - c. Select the **Minimum number of members that must be healthy for the group to be healthy** from the drop-down list. The default is **All**.
 - d. Add the health check members to the composite test from the **Available Aliases** list by selecting the health check to add and clicking **Add** to move the alias to the **Selected Alias** list.
 - e. To change the default notifications for this test, click **Override the default notifications**. By default, no notifications are sent for any health checks. Select the override options. You can cancel your choices by clicking **Clear all overrides**. For detailed information about configuring notifications, see "[Configuring Health Check Notifications](#)" on page 1369
 - f. Click **OK** to close the override dialog.
 - g. Click **OK** to close the edit dialog.


Click **Apply**. Copying and Deleting User-Defined Health Checks

Only user-defined health checks can be copied and deleted. Automatically generated health checks cannot be copied or deleted.

- ❑ If the source health check is user-defined host or a composite and the target alias name does not exist:
 - A new health check of the same kind with that alias name is created
 - The new health check has identical configuration settings to the source health check.
- ❑ If the target alias does exist and the target is of the same kind (that is, both are user-defined hosts or both are composite), then the complete configuration is copied from the source to the target.
- ❑ If a health check is referenced either in policy or in another health check, it cannot be deleted.

To copy or delete a user-defined host or composite health check:

1. Select **Configuration > Health Checks > General > Health Checks**.
2. Select the user-defined host or composite health check to copy or to delete.
3. Click **Copy** or **Delete**, as applicable.



If the target does not match the source type, the copy operation fails and you receive an error message.

Section I: Viewing Health Check Statistics

The topics in this section discuss health check statistics.

Health Check Topics

This section discusses the following topics:

- ❑ "Viewing Health Checks"
- ❑ "About Health Check Statistics" on page 1392
- ❑ "Interpreting Health Check Statistics" on page 1394

Viewing Health Checks

The ProxySG presents a comprehensive list of all the health checks configured on the appliance in the **Statistics > Health Checks** tab. You can view the details and events for each health check in this screen. To edit the health checks, go to the **Configuration > Health Checks > General** tab.

To view health checks on the ProxySG:

Select **Statistics > Health Checks**. The list of configured health checks displays.

Maintenance

Statistics

Health: Warning

Health Checks

Current Time: 2008-01-23 11:17:16

Last Boot: 2008-01-22 11:35:08

Since Boot: 23.7 hrs

Name ▲	Status	State	Last check		Duration	Since last transition			
			When	Time (ms)		#Checks	Avg (ms)	Min (ms)	Max (ms)
auth.bc_iwa	OK on alt server		11:17:12	6	17.0 min	20	8	5	11
auth.blue_coat_ldap	Unknown		-	-	-	-	-	-	-
dns.10.2.2.100	OK with errors		11:17:01	4	23.7 hrs	5,537,695	5	3	40
dns.10.2.2.101	Check failed		11:17:08	3	18.0 min	46	-	-	-
dns.172.16.90.110	OK		11:17:09	15	23.7 hrs	3,892,344,090	18	13	459
socks.gateway1	OK		11:16:49	12	3.7 hrs	2,837	65	18	1,614
drtr.rating_service	OK for some IPs		08:41:57	5,104	17.6 hrs	8,544	5,368	96	10,071
fwd.google	OK		11:17:00	28	20.2 hrs	10,790	47	28	3,072
fwd.my_ssh	OK		11:17:14	5	23.7 hrs	2,810	6	4	75
icap.inbound	OK		11:17:13	18	38.6 min	82	22	17	286
icap.outbound	OK		11:17:11	17	36.6 min	77	18	16	26
icap.test	Disabled: Healthy		1995-12-31	9,990	23.7 hrs	4,258	9,992	9,970	10,032
socks.personal	Disabled: Unhealthy		11:15:08	994	18.6 min	91	-	-	-
user.group	Check failed		1995-12-31	346	23.7 hrs	10,872	-	-	-
user.public.dns.server	OK		11:17:15	21	23.7 hrs	3,330,935,572	53	15	1,194

Details

Events

About Health Check Statistics

The **Statistics > Health Check** panel provides a snapshot of all the health checks configured on the device. By default, the screen is sorted by the name column. To change the sort order, click any column header to sort by that column.

The **Statistics > Health Check** screen displays the following information:

- ❑ **Current time:** Displays the current date and time.
- ❑ **Last Boot:** Displays the date and time when the device was last booted.
- ❑ **Since Boot:** Displays the time that the device has been functioning since the last boot.
- ❑ **Status:** Displays the summary of each health check configured on the ProxySG.
 - **Name:** The health check name. Example, `auth.blue_coat_iwa`
 - **State:** The health check state is represented by an icon and a status message. If the health check is disabled, it displays as:
 - Disabled: Healthy
 - Disabled: Unhealthy

If the health check is enabled, the table below shows the messages displayed:

Table 73–2 Status messages for enabled health checks

Status Message	Icon	Description	Health State
Unknown		Health has not yet been tested successfully.	Healthy
OK		The target device or service is completely healthy.	Healthy
OK with errors (multiple IP addresses)		One or more IP addresses have errors but none are down.	Healthy
OK for some IP addresses (multiple IP addresses)		One or more IP addresses are down but not all.	Healthy
OK on alt server		The primary server has failed; the realm is functioning on the alternate server.	Healthy
Functioning but going down (single IP address)		Failures are occurring; but the IP address is still functioning.	Healthy
Check failed		Device or service cannot be used.	Unhealthy
DNS failed		The hostname cannot be resolved	Unhealthy

- ❑ **Last check:** Information on the last completed health check probe.
 - **When:** Time of the last check.
 - **Time:** Response time of the last check.

- ❑ **Since last transition:** Displays aggregate values since the last transition between healthy and unhealthy.
 - **Duration:** Length of time since the last transition.
 - **#Checks:** Number of health checks performed since the last transition.
 - **Avg:** The mean response time since the last transition. This statistic is not displayed for a health check reporting unhealthy.
 - **Min:** Minimum response time. This statistic is not displayed for a health check reporting unhealthy.
 - **Max:** Maximum response time. This statistic is not displayed for a health check reporting unhealthy.
- ❑ **Details:** This option is active only if a single row is selected. When you click **Details**, it displays a new HTML window that contains detailed statistics on the selected health check. For example, in a domain check, this display provides an itemized explanation about each IP address in a domain.
- ❑ **Events:** This button is active only when a single row is selected. When you click the button, it displays a new HTML window containing the filtered event log entries for the selected health check.

Interpreting Health Check Statistics

The **Statistics > Health Check** tab in the Management Console provides a snapshot of all the health checks configured on the ProxySG. This screen allows you to glance at the health checks for routine maintenance, to diagnose potential problems, and to view health check failures.

The following figure shows the **Statistics > Health Check** panel along with an explanation of the display.

Maintenance

Statistics

Health: Warning

Health Checks

Current Time: 2008-01-23 11:17:16

Last Boot: 2008-01-22 11:35:08

Since Boot: 23.7 hrs

Name ▲	Status	Last check		Since last transition					
		When	Time (ms)	Duration	#Checks	Avg (ms)	Min (ms)	Max (ms)	
auth.bc_iwa	OK on alt server	11:17:12	6	17.0 min	20	8	5	11	
auth.blue_coat_ldap	Unknown	-	-	-	-	-	-	-	
dns.10.2.2.100	OK with errors	11:17:01	4	23.7 hrs	5,537,695	5	3	40	
dns.10.2.2.101	Check failed	11:17:08	3	18.0 min	46	-	-	-	
dns.172.16.90.110	OK	11:17:09	15	23.7 hrs	3,892,344,090	18	13	459	
socks.gateway1	OK	11:16:49	12	3.7 hrs	2,837	65	18	1,614	
drtr.rating_service	OK for some IPs	08:41:57	5,104	17.6 hrs	8,544	5,368	96	10,071	
fwd.google	OK	11:17:00	28	20.2 hrs	10,790	47	28	3,072	
fwd.my_ssh	OK	11:17:14	5	23.7 hrs	2,810	6	4	75	
icap.inbound	OK	11:17:13	18	38.6 min	82	22	17	286	
icap.outbound	OK	11:17:11	17	36.6 min	77	18	16	26	
icap.test	Disabled: Healthy	1995-12-31	9,990	23.7 hrs	4,258	9,992	9,970	10,032	
socks.personal	Disabled: Unhealthy	11:15:08	994	18.6 min	91	-	-	-	
user.group	Check failed	1995-12-31	346	23.7 hrs	10,872	-	-	-	
user.public.dns.server	OK	11:17:15	21	23.7 hrs	3,330,935,572	53	15	1,194	

Details

Events

- ❑ The current time is 11:17 on January 23, 2008
- ❑ Authentication realm Symantec IWA —auth.bc_iwa is functioning on its alternate server for 17 minutes. The primary server failed just 17 minutes ago.
- ❑ Authentication realm Symantec LDAP — auth.blue_coat_ldap is configured, but is not currently referenced in policy. The health state is **Unknown** because it is not being queried by the ProxySG for authentication lookups or for health checks.
- ❑ DNS server —dns.10.2.2.100 is functioning with errors, and it reports healthy since boot. Select the row and click **Events** to view the expanded display about the earlier failed health check.
- ❑ DNS server —dns.10.2.2.101 is not functioning since 11:17 (for 18 minutes now). Select the row and click **Details** to view the expanded display for the health check.
- ❑ DNS server 172.16.90.110 reports healthy and is stable since the device was booted.
- ❑ SOCKS gateway— socks.gateway1 is healthy and is operating for the last 3.7 hours. The average response time for this gateway is 65 ms.
- ❑ WebPulse service group — WebPulse.rating-service is healthy, and the average response time is adequate. However the status icon shows that the service is experiencing difficulties with some IP addresses. Select the row and click **Details** to view the information on the configured IP addresses and the failure points. The **Details** button displays the following information:

Domain name: sp.cwfservice.net DNS status: success

```
Enabled OK for some IPs UP
IP address: 217.169.46.101 Enabled OK UP
  Last status: Success.
  Successes (total): 8 (last): Wed, 23 Jan 2008 16:35:36 GMT
  (consecutive): 8
  Failures (total): 0 (last): Never (consecutive): 0 (external): 0
  Last response time: 331 ms Average response time: 357 ms
  Minimum response time: 300 ms Maximum response time: 613 ms
IP address: 65.160.238.181 Enabled Check failed DOWN
  Last status: A communication error has occurred.
  Successes (total): 0 (last): Never (consecutive): 0
  Failures (total): 3809 (last): Wed, 23 Jan 2008 16:45:09 GMT
  (consecutive): 3809 (external): 0
  Last response time: 9990 ms Average response time: 9992 ms
  Minimum response time: 9981 ms Maximum response time: 10071 ms
IP address: 204.246.129.201 Enabled OK UP
  Last status: Success.
  Successes (total): 8 (last): Wed, 23 Jan 2008 16:41:57 GMT
  (consecutive): 6
  Failures (total): 15 (last): Wed, 23 Jan 2008 01:41:44 GMT
  (consecutive): 0 (external): 0
  Last response time: 104 ms Average response time: 1133 ms
  Minimum response time: 96 ms Maximum response time: 6281 ms
IP address: 65.160.238.183 Enabled Check failed DOWN
  Last status: A communication error has occurred.
  Successes (total): 0 (last): Never (consecutive): 0
  Failures (total): 3809 (last): Wed, 23 Jan 2008 16:45:09 GMT
  (consecutive): 3809 (external): 0
  Last response time: 9991 ms Average response time: 9993 ms
  Minimum response time: 9981 ms Maximum response time: 10067 ms
```

- ❑ Forwarding host — fwd.google is functioning for 20.2 hours.
- ❑ The forwarding host — fwd.my_ssh is healthy since boot.
- ❑ ICAP service — icap.inbound and icap.outbound are healthy.
- ❑ ICAP service — icap.test is disabled and configured to report healthy. Disabled health checks appear grayed out on the screen.
- ❑ SOCKS gateway — socks.personal is disabled and configured to report unhealthy.
- ❑ The user-defined health check — user.public.dns.server is healthy.

Section J: Using Policy

The results of a health check can be affected through forwarding, SOCKS gateway, or SSL certificate policy. The health check transactions execute the `<forward>` layer and (for SSL or HTTPS tests) the `<ssl>` layer to determine applicable policy.

This allows health check behavior to match as closely as possible to that of the SSL traffic that the health check is monitoring.

Health checks cannot be deleted while referenced in policy. If a health check is automatically deleted when its target is deleted, a reference to the health check in policy can block deletion not only of the health check but of its target.

Two policy conditions exist for health checks:

- ❑ `health_check=` : This condition tests whether the current transaction is a health check transaction. Optionally, the condition tests whether the transaction is that of a specific health check.
- ❑ `is_healthy.health_check_name=` : This condition tests whether the specified health check is healthy.

Example: For a user-defined health check `user.internet` that gates access to a popular Web site and tests for Internet connectivity and responsiveness, you could define policy to redirect traffic through a forwarding host if the health check fails.

To do this in policy:

```
<Forward>
  is_healthy.user.internet=no forward(alternate_route)
```

For more information about using policy, refer to the *Visual Policy Manager Reference* and *Content Policy Language Guide*.

Chapter 74: Maintaining the ProxySG

The following sections describe how to maintain the ProxySG. It includes the following topics:

- ❑ "Restarting the ProxySG" on page 1399
- ❑ "Restoring System Defaults" on page 1401
- ❑ "Clearing the DNS Cache" on page 1403
- ❑ "Clearing the Object Cache" on page 1403
- ❑ "Clearing the Byte Cache" on page 1404
- ❑ "Clearing Trend Statistics" on page 1404
- ❑ "Upgrading the ProxySG Appliance" on page 1404
- ❑ "Managing ProxySG Systems" on page 1404
- ❑ "Disk Reinitialization" on page 1407
- ❑ "Deleting Objects from the ProxySG Appliance" on page 1408

Performing Maintenance Tasks

You can perform the following maintenance tasks on the ProxySG:


- ❑ "Restarting the ProxySG" on page 1399
- ❑ "Shutting Down the ProxySG VA" on page 1401
- ❑ "Restoring System Defaults" on page 1401
- ❑ "Clearing the DNS Cache" on page 1403
- ❑ "Clearing the Object Cache" on page 1403
- ❑ "Clearing the Byte Cache" on page 1404
- ❑ "Clearing Trend Statistics" on page 1404

Restarting the ProxySG

When you restart the ProxySG, you can choose between a software only restart or a hardware and software restart as follows.

To restart the ProxySG appliance:

1. Select **Maintenance > System and Disks > Tasks**.



2. In the **Maintenance Tasks** field, select one of the following options:
 - **Software Only**—Applicable for most situations, such as suspected system hang.
 - **Hardware and software**—A more comprehensive restart, this option might take several minutes longer depending on the amount of memory and the number of disk drives present. Symantec recommends this option if a hardware fault is suspected.
3. (Hardware and software restart only) Select a system that you want to start upon reboot from the **System to run** drop-down list (the default system is pre-selected).
4. (Optional) Click **Apply** if you want the restart options to be the default upon the next system restart.
5. Click **Restart now**. The Restart System dialog displays.
6. To proceed with the restart, click **OK**.

See Also

["Restoring System Defaults" on page 1401](#)

["Restore-Defaults" on page 1401](#)

["Clearing the DNS Cache" on page 1403](#)

["Clearing the Object Cache" on page 1403](#)

["Clearing the Byte Cache" on page 1404](#)

["Clearing Trend Statistics" on page 1404](#)

Related CLI Syntax to Configure the Hardware/Software Restart Settings

```
SGOS#(config) restart mode {hardware | software}
SGOS# restart abrupt
SGOS# restart regular
SGOS# restart upgrade
```


Shutting Down the ProxySG VA

If you need to reboot the ProxySG VA, you should “gracefully” shut it down using the procedure below. You should shut down the system before performing the following tasks: system backup, server software upgrade, taking the server offline for maintenance, migration of the ProxySG VA to a different server, installing additional or higher-capacity drives on the ESX host, and ProxySG VA configuration (for example, adding a serial port, upgrading the model, and so forth).

To shut down the ProxySG VA:

1. Select **Maintenance > System and Disks > Tasks**.
2. In the **Maintenance Tasks** area, click **Shut down**. The System Shutdown dialog box displays.
3. Click **OK**.

Your VMware client also offers a command for powering down a virtual machine. This is an alternate way to shut down the ProxySG VA.

Restoring System Defaults

SGOS allows you to restore some or all of the system defaults. Use these commands with caution. The `restore-defaults` command deletes most, but not all, system defaults:

- ❑ The `restore-defaults` command with the `factory-defaults` option reinitializes the ProxySG to the original settings it had when it was shipped from the factory. You must use the CLI to perform this action.
- ❑ The `restore-defaults` command with the `keep-console` option restores the default settings without losing all IP addresses on the system. This action is available in the Management Console and the CLI.

The following sections describe the three possible operations:

- ❑ ["Restore-Defaults" on page 1401](#)
- ❑ ["Keep-Console" on page 1402](#)
- ❑ ["Factory-Defaults" on page 1403](#)

Restore-Defaults

Settings that are deleted when you use the `restore-defaults` command include:

- ❑ All IP addresses (these must be restored before you can access the Management Console again).
- ❑ DNS server addresses (these must be restored through the CLI before you can access the Management Console again).
- ❑ Installable lists.
- ❑ All customized configurations.
- ❑ Blue Coat trusted certificates.

- ❑ Original SSH (v1 and v2) host keys (new host keys are regenerated).

You can use the `force` option to restore defaults without confirmation.

Keep-Console

Settings that are retained when you use the `restore-defaults` command with the `keep-console` option include:

- ❑ IP interface settings, including VLAN configuration.
- ❑ Default gateway and static routing configuration.
- ❑ Virtual IP address configuration.
- ❑ Bridging settings.
- ❑ Failover group settings.

Using the `keep-console` option retains the settings for all consoles (Telnet, SSH, HTTP, and HTTPS), whether they are enabled, disabled, or deleted.


Administrative access settings retained using the `restore-defaults` command with the `keep-console` option include:

- ❑ Console username and password.
- ❑ Front panel pin number.
- ❑ Console enable password.
- ❑ SSH (v1 and v2) host keys.
- ❑ Keyrings used by secure console services.
- ❑ RIP configurations.

You can also use the `force` option to restore defaults without confirmation.

To perform a `restore-defaults keep-console` action using the Management Console:

1. Select the **Maintenance > System and Disks > Tasks** tab.



2. In the **Maintenance Tasks** field, click **Restore**. This invokes the `restore-defaults keep-console` action. The Restore Configuration dialog displays.
3. Click **OK**. The following settings are retained:
 - IP addresses, including default gateway and bridging (virtual IP addresses are *not* retained).

- Settings for all consoles.
- Ethernet maximum transmission unit (MTU) size.
- TCP round trip time.
- Static routes table information.

To perform a restore-defaults keep-console action using the CLI:

Enter the following command:

```
SGOS# restore-defaults keep-console
```

Factory-Defaults

All system settings are deleted when you use the `restore-defaults` command with the `factory-defaults` option.

The only settings that are retained are:

- ❑ Trial period information
- ❑ The last five installed appliance systems, from which you can pick one for rebooting

The Serial Console password is also deleted if you use `restore-defaults factory-defaults`. For information on the Serial Console password, see "[Securing the Serial Port](#)" on page 60.

You can use the `force` option to restore defaults without confirmation.

To restore the system to the factory defaults using the CLI:

Enter the following command:

```
SGOS# restore-defaults factory-defaults
```

Clearing the DNS Cache

You can clear the DNS cache at any time. You might need to do so if you have experienced a problem with your DNS server or if you have changed your DNS configuration.

To clear the DNS cache:

1. Select the **Maintenance > System and disks > Tasks** tab.
2. In the **Cache and Statistics Tasks** field, click **Clear** next to **the DNS cache**. The Clear System DNS Cache dialog displays.
3. Click **OK**.

Clearing the Object Cache

You can clear the object cache at any time.

When you clear the cache, all objects in the cache are set to *expired*. The objects are not immediately removed from memory or disk, but a subsequent request for any object requested is retrieved from the source before it is served.

To clear the object cache:

1. Select the **Maintenance > System and disks > Tasks** tab.
2. In the **Cache and Statistics Tasks** field, click **Clear** next to the **object cache**. The Clear Object Cache dialog displays.
3. Click **OK**.

Clearing the Byte Cache

You can clear the byte cache at any time. A user case to perform this action is testing purposes.

To clear the byte cache:

1. Select the **Maintenance > System and disks > Tasks** tab.
2. In the **Cache and Statistics Tasks** field, click **Clear** next to the **byte cache**. The Clear Byte Cache dialog displays.
3. Click **OK**.

Clearing Trend Statistics

You can clear all trend statistics at any time.

To clear all trend statistics:

1. Select the **Maintenance > System and disks > Tasks** tab.
2. In the **Cache and Statistics Tasks** field, click **Clear** next to the **trend statistics**. The Clear Trend Statistics dialog displays.
3. Click **OK**.

Upgrading the ProxySG Appliance

Before upgrading the appliance, refer to the *SGOS Upgrade/Downgrade Quick Reference* to determine your upgrade path:

<http://www.symantec.com/docs/DOC9794>

Once you have determined your upgrade path, refer to the *SGOS Upgrade/Downgrade Guide* to upgrade the appliance.

Managing ProxySG Systems

The ProxySG **Systems** tab displays the five available systems. Empty systems are indicated by the word **Empty**.

The system currently running is highlighted in blue and cannot be replaced or deleted.


From this screen, you can:

- ❑ View details of the available SGOS system versions.

- ❑ Select the SGOS system version to boot. See ["Setting the Default Boot System"](#) on page 1406.
- ❑ Lock one or more of the available SGOS system versions. See ["Locking and Unlocking ProxySG Systems"](#) on page 1406.
- ❑ Select the SGOS system version to be replaced. See ["Replacing a ProxySG System"](#) on page 1407.
- ❑ Delete one or more of the available SGOS system versions (CLI only). See ["Deleting a ProxySG System"](#) on page 1407.


To view SGOS system replacement options:

Select the **Maintenance > Upgrade > Systems** tab.



To view details for an SGOS system version:

1. Select the **Maintenance > Upgrade > Systems** tab.
2. Click **Details** next to the system for which you want to view detailed information; click **OK** when you are finished.



To view details for an SGOS system version:

At the command prompt:

```
SGOS> show installed-systems
```

Example Session

```
SGOS> show installed-systems
ProxySG Appliance Systems
1. Version: SGOS 5.4.1.3, Release ID: 25460
   Thursday June 25 2009 08:49:55 UTC, Lock Status: Locked
   Boot Status: Last boot succeeded, Last Successful Boot: Thursday
   April 6 2006 17:33:19 UTC
2. Version: SGOS 5.4.1.1, Release ID: 25552 Debug
   Friday April 14 2009 08:56:55 UTC, Lock Status: Unlocked
   Boot Status: Last boot succeeded, Last Successful Boot: Friday April
   14 2006 16:57:18 UTC
3. Version: N/A, Release ID: N/A ( EMPTY )
   No Timestamp, Lock Status: Unlocked
   Boot Status: Unknown, Last Successful Boot: Unknown
4. Version: N/A, Release ID: N/A ( EMPTY )
   No Timestamp, Lock Status: Unlocked
   Boot Status: Unknown, Last Successful Boot: Unknown
5. Version: N/A, Release ID: N/A ( EMPTY )
   No Timestamp, Lock Status: Unlocked
   Boot Status: Unknown, Last Successful Boot: Unknown
   Default system to run on next hardware restart: 2
   Default replacement being used. (oldest unlocked system)
   Current running system: 2

When a new system is loaded, only the system number that was replaced
is changed.

The ordering of the rest of the systems remains unchanged.
```

Setting the Default Boot System

This setting allows you to select the system to be booted on the next hardware restart. If a system starts successfully, it is set as the default boot system. If a system fails to boot, the next most recent system that booted successfully becomes the default boot system.

To set the ProxySG appliance to run on the next hardware restart:

1. Select the **Maintenance > Upgrade > Systems** tab.
2. Select the preferred System version in the **Default** column.
3. Click **Apply**.

Note: An empty system cannot be specified as default, and only one system can be specified as the default system.

Locking and Unlocking ProxySG Systems

Any system can be locked, except a system that has been selected for replacement. If all systems, or all systems except the current system, are locked, the ProxySG cannot load a new system.

If a system is locked, it cannot be replaced or deleted.

To lock a system:

1. Select the **Maintenance > Upgrade > Systems** tab.
2. Select the system(s) to lock in the **Lock** column.
3. Click **Apply**.

To unlock a system:

1. Select the **Maintenance > Upgrade > Systems** tab.
2. Deselect the system(s) to unlock in the **Lock** column.
3. Click **Apply**.

Replacing a ProxySG System

You can specify the system to be replaced when a new system is downloaded. If no system is specified, the oldest unlocked system is replaced by default. You cannot specify a locked system for replacement.

To specify the system to replace:

1. Select the **Maintenance > Upgrade > Systems** tab.
2. Select the system to replace in the **Replace** column.
3. Click **Apply**.

Deleting a ProxySG System

You can delete any of the system versions except the current running system. A locked system must be unlocked before it can be deleted. If the system you want to delete is the default boot system, you need to select a new default boot system before the system can be deleted.

You cannot delete a system version through the Management Console; you must use the CLI.

To delete a system:

At the (config) command prompt:

```
SGOS#(config) installed-systems  
SGOS#(config installed-systems) delete system_number
```

where *system_number* is the system you want to delete.

Disk Reinitialization

You can reinitialize disks on a multi-disk ProxySG. You cannot reinitialize the disk on a single-disk ProxySG. If you suspect a disk fault in a single-disk system, contact Blue Coat Technical Support for assistance.

About Reinitialization

Reinitialization is done online without rebooting the system. (For more information, refer to the #disk command in the *Command Line Interface Reference*.)

Important: Do not reinitialize disks while the system is proxying traffic.

SGOS operations, in turn, are not affected, although during the time the disk is being reinitialized, that disk is not available for caching. Only the master disk reinitialization restarts the ProxySG.

Only persistent objects are copied to a newly-reinitialized disk. This is usually not a problem because most of these objects are replicated or mirrored. If the reinitialized disk contained one copy of these objects (which is lost), another disk contains another copy.

You cannot reinitialize all of the ProxySG disks over a very short period of time. Attempting to reinitialize the last disk in a system before critical components can be replicated to other disks in the system causes a warning message to appear.

Immediately after reinitialization is complete, the ProxySG automatically starts using the reinitialized disk for caching.

Note: If a disk containing an unmirrored event or access log is reinitialized, the logs are lost. Similarly, if two disks containing mirrored copies of the logs are reinitialized, both copies of the logs are lost.

Hot Swapping Disk Drives in 810 and 8100 ProxySG Appliances

On multi-disk 810 and 8100 ProxySG appliances, you can hot swap any disk (including the left-most disk, which on earlier appliances was known as the master disk—the newer platforms do not have this concept) as long as there is one operational disk drive. When you hot swap a disk drive, the data on the existing disk is transferred to the new disk and vice versa. Because the data from each disk is copied back and forth, you might need to change the default boot version. This is because the ProxySG always boots the newest OS—if the disk drive had a newer OS, the ProxySG tries to boot it—even if you had previously set a different default boot version. Thus, you should reset your default boot version after hot swapping a disk drive. See "[Setting the Default Boot System](#)" on page 1406 for more information.

Single-Disk ProxySG Appliance

The disk on a single-disk ProxySG cannot be reinitialized by the customer. If you suspect a disk fault in a single-disk ProxySG, contact Blue Coat Technical Support for assistance.

Deleting Objects from the ProxySG Appliance

The ability to delete either individual or multiple objects from the ProxySG makes it easy to delete stale or unused data and make the best use of the storage in your system.

Note: The maximum number of objects that can be stored in a ProxySG is affected by a number of factors, including the SGOS version it is running and the hardware platform series.

This feature is not available in the Management Console. Use the CLI instead.

To delete a single object from the ProxySG:

At the (config) prompt, enter the following command:

```
SGOS#(config) content delete url url
```

To delete multiple objects from the ProxySG:

At the (config) prompt, enter the following command:

```
SGOS#(config) content delete regex regex
```


Chapter 75: Diagnostics

This chapter describes the various resources that provide diagnostic information.

Topics in this Chapter

This chapter includes information about the following topics:

- ❑ "Diagnostic Terminology"
- ❑ "Diagnostic Reporting (Service Information)" on page 1412 (This includes taking snapshots of the system.)
- ❑ "Packet Capturing (PCAP—the Job Utility)" on page 1419
- ❑ "Core Image Restart Options" on page 1425
- ❑ "Diagnostics: Symantec Customer Experience Program and Monitoring" on page 1426
- ❑ "Diagnostic Reporting (CPU Monitoring)" on page 1427

If the ProxySG does not appear to work correctly and you are unable to diagnose the problem, contact Blue Coat Technical Support.

Diagnostic Terminology

- ❑ Heartbeats: Enabled by default, Heartbeats (statistics) are a diagnostic tool used by Blue Coat, allowing them to proactively monitor the health of appliances.
- ❑ Core images: Created when there is an unexpected system restart. This stores the system state at the time of the restart, enhancing the ability for Blue Coat to determine the root cause of the restart.
- ❑ SysInfo (System Information): SysInfo provides a snapshot of statistics and events on the ProxySG.
- ❑ PCAP: An onboard packet capture utility that captures packets of Ethernet frames going in or out of an ProxySG.
- ❑ Policy trace: A policy trace can provide debugging information on policy transactions. This is helpful, even when policy is not the issue. For information on using policy tracing, refer to the *Content Policy Language Guide*.
- ❑ Policy coverage: This feature reports on the rules and objects that match user requests processed through the appliance's current policy. For more information on policy coverage, refer to the "Troubleshooting" chapter in the *Content Policy Language Guide* and the following article:

<http://www.symantec.com/docs/TECH241425>

- ❑ **Event Logging:** The event log files contain messages generated by software or hardware events encountered by the appliance. For information on configuring event logging, see ["Configuring Event Logging and Notification"](#) on page 1309.
- ❑ **Access Logging:** Access logs allow for analysis of Quality of Service, content retrieved, and other troubleshooting. For information on Access Logging, see ["About Access Logging"](#) on page 617.
- ❑ **CPU Monitoring:** With CPU monitoring enabled, you can determine what types of functions are taking up the majority of the CPU.

To test connectivity, use the following commands from the enable prompt:

- ❑ `ping`: Verifies that a particular IP address exists and is responding to requests.
- ❑ `traceroute`: Traces the route from the current host to the specified destination host.
- ❑ `test http get path_to_URL`: Makes a request through the same code paths as a proxied client.
- ❑ `display path_to_URL`: Makes a direct request (bypassing the cache).
- ❑ `show services`: Verifies the port of the Management Console configuration.
- ❑ `show policy`: Verifies if policy is controlling the Management Console.

For information on using these commands, refer to Chapter 2: "Standard and Privileged Mode Commands" in the *Command Line Interface Reference*.

Note: If you cannot access the Management Console at all, ensure that you are using HTTPS (`https://ProxySG_IP_address:8082`). To use HTTP, you must explicitly enable it before you can access the Management Console.

Diagnostic Reporting (Service Information)

The service information options allow you to send service information to Blue Coat using either the Management Console or the CLI. You can select the information to send, send the information, view the status of current transactions, and cancel current transactions. You can also send service information automatically in case of a crash.

Sending Service Information Automatically


Enabling automatic service information allows you to enable the transfer of relevant service information automatically whenever a crash occurs. This saves you from initiating the transfer, and increases the amount of service information that Blue Coat can use to solve the problem. The core image, system configuration, and event log are system-use statistics that are sent for analysis. If a packet capture exists, it is also sent.

The auto-send feature requires that a valid Service Request is entered. If you do not have a Service Request open you must first contact Blue Coat Technical Support.

Important: A core image and packet capture can contain sensitive information—for example, parts of an HTTP request or response. The transfer to Blue Coat is encrypted, and therefore secure; however, if you do not want potentially sensitive information to be sent to Blue Coat automatically, do not enable the automatic service information feature.

To send service information automatically:

1. Select the **Maintenance > Service Information > Send Information > General** tab.



2. To send core image service information to Blue Coat automatically, select **Enable auto-send**.
3. Enter the service-request number that you received from a Technical Support representative into the **Auto Send Service Request Number** field (the service-request number is in the form `xx-xxxxxxx` or `x-xxxxxxx`).
4. Click **Apply**.
5. (Optional) To clear the service-request number, clear the **Auto Send Service Request Number** field and click **Apply**.

Managing the Bandwidth for Service Information

You can control the allocation of available bandwidth for sending service information. Some service information items are large, and you might want to limit the bandwidth used by the transfer. Changing to a new bandwidth management class does not affect service information transfers already in progress. However, changing the details of the bandwidth management class used for service information, such as changing the minimum or maximum bandwidth settings, affects transfers already in progress if that class was selected prior to initiating the transfer.

Note: Before you can manage the bandwidth for the automatic service information feature, you must first create an appropriate bandwidth-management class. For information about creating and configuring bandwidth classes, see "[Configuring Bandwidth Allocation](#)" on page 602.

To manage bandwidth for service information:

1. Select the **Maintenance > Service Information > Send Information > General** tab.
2. To manage the bandwidth of automatic service information, select a bandwidth class from the **Service Information Bandwidth Class** drop-down menu.

3. Click **Apply**.
4. (Optional) To disable the bandwidth-management of service information, select **none** from the **Service Information Bandwidth Class** drop-down menu; click **Apply**.

Configure Service Information Settings

The service information options allow you to send service information to Blue Coat using either the Management Console or the CLI. You can select the information to send, send the information, view the status of current transactions, and cancel current transactions using either the Management Console or the CLI. For information about sending service information automatically, see “[Sending Service Information Automatically](#)” on page 82.

Important: You must specify a service-request number before you can send service information. See Blue Coat Technical Support at: <http://www.bluecoat.com/support> for details on opening a service request ticket.

The following list details information that you can send:

- ☐ Packet Capture
- ☐ Event Log
- ☐ Memory Core
- ☐ Policy Trace File
- ☐ SYSInfo
- ☐ Access Logs (can specify multiple)
- ☐ Snapshots (can specify multiple)
- ☐ Contexts (can specify multiple)

To send service information:


1. Select the **Maintenance > Service Information > Send Information > Send Service Information** tab.

The screenshot shows the 'Send Service Information' dialog box. The 'General' tab is selected. The 'Service Request Number' field contains '12-34567879'. The 'Information to send' section has checkboxes for 'Packet Capture (Unknown)', 'Memory Core (Unknown)', 'Policy Trace File (Unknown)', 'Access Logs', 'Snapshots', and 'Contexts'. The 'Snapshots' checkbox is checked. To the right of these checkboxes are three buttons: 'Select access logs to send', 'Select snapshots to send', and 'Select contexts to send'. The 'Select snapshots to send' button is highlighted with an orange border. At the bottom left is a 'Select Newest' button, and at the bottom right are 'Send' and 'View Progress' buttons. Arrows labeled 2a, 2b, and 2c point to the 'Service Request Number' field, the 'Information to send' section, and the 'Snapshots' checkbox/button respectively.


2. Select options as required:
 - a. Enter the service-request number that you received from a Technical Support representative. The service-request number format is:
x-xxxxxxxxx
 - b. Select the appropriate options (as indicated by a Technical Support representative) in the **Information to send** area.

Note: Options for items that you do not have on your system are grayed out and cannot be selected.

- c. (Optional) If you select **Access Logs**, **Snapshots**, or **Contexts**, you must also click **Select access logs to send**, **Select snapshots to send**, or **Select contexts to send** and complete the following steps in the corresponding dialog that displays:



- d. To select information to send, highlight the appropriate selection in the **Access Logs/Snapshots/Contexts Not Selected** field and click **Add to Selected**.
 - e. To remove information from the **Access Logs/Snapshots/Contexts Selected** field, highlight the appropriate selection and click **Remove from Selected**.
 - f. Click **Ok** to close the dialog.
3. Click **Send**.



4. Click **Ok** in the Information upload started dialog that appears.

Creating and Editing Snapshot Jobs

The snapshot subsystem periodically pulls a specified console URL and stores it in a repository, offering valuable resources for Blue Coat customer support in diagnosing problems.

By default, two snapshots are defined:

- ❑ **sysinfo:** Takes a snapshot of the system information URL once every 24 hours. This snapshot job keeps the last 100 snapshots.
- ❑ **sysinfo_stats:** Takes an hourly snapshot of the system information statistics (sysinfo_stats). This snapshot job keeps the last 168 snapshots.

Determining which console URL to poll, the time period between snapshots, and how many snapshots to keep are all configurable options for each snapshot job.

Compatibility With Pre-6.5.2 Snapshots


Changes were made to the snapshots in SGOS 6.5.2. Note the following:

- ❑ Snapshots created in SGOS 6.5.2 or later are not viewable if you downgrade to SGOS 6.5.1 or earlier.
- ❑ When running SGOS 6.5.2, you can view snapshots taken by a previous SGOS version at the following URL:

`/Diagnostic/Snapshot/Old`

To create a new snapshot job:

1. Select the **Maintenance > Service Information > Snapshots** tab.




2. Perform the following steps:
 - a. Click **New**.

- b. Enter a snapshot job into the Add list item dialog that displays; click **Ok**.
3. Click **Apply**.
4. (Optional) To view snapshot job information, click **View All Snapshots**. Close the window that opens when you are finished viewing.

To edit an existing snapshot job:

1. Select **Maintenance > Service Information > Snapshots**.
2. Select the snapshot job you want to edit (highlight it).
3. Click **Edit**.

The Edit Snapshot dialog displays.



4. Enter the following information into the Edit Snapshot fields:
 - a. **Target:** Enter the object to snapshot.
 - b. **Interval (minutes):** Enter the interval between snapshot reports.
 - c. **Total Number To Take:** Enter the total number of snapshots to take or select **Infinite** to take an infinite number of snapshots.
 - d. **Maximum Number To Store:** Enter the maximum number of snapshots to store. The maximum number of snapshots you can store is now 1000 (it was 100 in previous versions).
 - e. **Enabled:** Select this to enable this snapshot job or deselect it to disable this snapshot job.

5. (Optional) Click **View URL List** to open a window displaying a list of URLs; close the window when you are finished viewing.
6. (Optional) Click **View Snapshots** to open a window displaying snapshot information; close the window when you are finished viewing.
7. (Optional) Click **Clear Snapshots** to clear all stored snapshot reports.

Packet Capturing (PCAP—the Job Utility)

You can capture packets of Ethernet frames going into or leaving a ProxySG. Packet capturing allows filtering on various attributes of the frame to limit the amount of data collected. Any packet filters must be defined before a capture is initiated, and the current packet filter can only be modified if no capture is in progress.

The `pcap` utility captures all received packets that are either directly addressed to the ProxySG through an interface's MAC address or through an interface's broadcast address. The utility also captures transmitted packets that are sent from the appliance. The collected data can then be transferred to the desktop or to Blue Coat for analysis.

Note: Packet capturing increases the amount of processor usage performed in TCP/IP.

To analyze captured packet data, you must have a tool that reads Packet Sniffer Pro 1.1 files (for example, Wireshark or Packet Sniffer Pro 3.0).

PCAP File Size

The PCAP file size is limited to 3% of the available system memory at startup (not to exceed 4GB). The default packet capture file size is 100MB.


The file size can be changed by specifying a value for the following options in the **Maintenance > Service Information > Packet Captures > Start Capture** dialog:

- ☐ **Capture first *n* matching Kilobytes**
- ☐ **Capture last *n* matching Kilobytes**

If both values are both specified, the maximum of the two values is used. See ["Configuring Packet Capturing"](#) on page 1421.

Determine Maximum File Size

To determine the maximum PCAP file size for your appliance, enter the value **9999999** into the **Capture first *n* matching Kilobytes** field and click **Start Capture**. The capture will terminate; the valid values are reported in red text.



PCAP File Name Format

The name of a downloaded packet capture file has the format:

`bluecoat_date_filter-expression.cap`, revealing the date and time (UTC) of the packet capture and any filter expressions used. Because the filter expression can contain characters that are not supported by a file system, a translation can occur. The following characters are not translated:

- ❑ Alphanumeric characters (a-z, A-Z, 0-9)
- ❑ Periods (.)

Characters that are translated are:

- ❑ Space (replaced by an underscore)
- ❑ All other characters (including the underscore and dash) are replaced by a dash followed by the ASCII equivalent; for example, a dash is translated to `-2D` and an ampersand (&) to `-26`.

Common PCAP Filter Expressions

Packet capturing allows filtering on various attributes of the frame to limit the amount of data collected. PCAP filter expressions can be defined in the Management Console or the CLI. Below are examples of filter expressions; for PCAP configuration instructions, see ["Configuring Packet Capturing"](#) on page 1421.

Some common filter expressions for the Management Console and CLI are listed below. The filter uses the Berkeley Packet Filter format (BPF), which is also used by the `tcpdump` program. A few simple examples are provided below. If filters with greater complexity are required, you can find many resources on the Internet and in books that describe the BPF filter syntax.

Note: Some qualifiers must be escaped with a backslash because their identifiers are also keywords within the filter expression parser.

- ❑ `ip proto protocol`
where *protocol* is a number or name (`icmp`, `udp`, `tcp`).
 - ❑ `ether proto protocol`
where *protocol* can be a number or name (`ip`, `arp`, `rarp`).
-

Table 75–1 PCAP Filter Expressions

Filter Expression	Packets Captured
<code>ip host 10.25.36.47</code>	Captures packets from a specific host with IP address 10.25.36.47.
<code>not ip host 10.25.36.47</code>	Captures packets from all IP addresses except 10.25.36.47.
<code>ip host 10.25.36.47 and ip host 10.25.36.48</code>	Captures packets sent between two IP addresses: 10.25.36.47 and 10.25.36.48. Packets sent from one of these addresses to other IP addresses are not filtered.
<code>ether host 00:e0:81:01:f8:fc</code>	Captures packets to or from MAC address 00:e0:81:01:f8:fc:.
<code>port 80</code>	Captures packets to or from port 80.
<code>ip sr www.bluecoat.com and ether broadcast</code>	Captures packets that have IP source of www.bluecoat.com and ethernet broadcast destination.

Using Filter Expressions in the CLI

To add a filter to the CLI, use the command:

```
SGOS# pcap filter expr parameters
```

To remove a filter, use the command:

```
SGOS# pcap filter <enter>
```

Important: Define CLI filter expr parameters with double-quotes to avoid confusion with special characters. For example, a space is interpreted by the CLI as an additional parameter, but the CLI accepts only one parameter for the filter expression. Enclosing the entire filter expression in quotations allows multiple spaces in the filter expression.

Configuring Packet Capturing


Use the following procedures to configure packet capturing. If a download of the captured packets is requested, packet capturing is implicitly stopped. In addition to starting and stopping packet capture, a filter expression can be configured to control which packets are captured. For information on configuring a PCAP filter, see "[Common PCAP Filter Expressions](#)" on page 1420.

Note: Requesting a packet capture download stops packet capturing.

To analyze captured packet data, you must have a tool that reads Packet Sniffer Pro 1.1 files (for example, Ethereal or Packet Sniffer Pro 3.0).


To enable, stop, and download packet captures:

1. Select the **Maintenance > Service Information > Packet Captures** tab.



2. Perform the following steps:

- a. In the **Direction** drop-down list, select the capture direction: **in**, **out**, or **both**.
- b. In the **Interface** drop-down list, select the interface on which to capture.
- c. To define or change the PCAP filter expression, enter the filter information into the **Capture filter** field. (See "[Common PCAP Filter Expressions](#)" on page 1420 for information about PCAP filter expressions for this field.) To remove the filter, clear this field.
- d. Click **Start Capture**. The Start Capture dialog displays.




3. Select options, as required:

- a. Select a buffer size:
 - Capture all matching packets.

- Capture first n matching packets. Enter the number of matching packets (n) to capture. If the number of packets reaches this limit, packet capturing stops automatically. The value must be between 1 and 1000000.
 - Capture last n matching packets. Enter the number of matching packets (n) to capture. Any packet received after the memory limit is reached results in the discarding of the oldest saved packet prior to saving the new packet. The saved packets in memory are written to disk when the capture is stopped. The value must be between 1 and 1000000.
 - Capture first n matching Kilobytes. Enter the number of kilobytes (n) to capture. If the buffer reaches this limit, packet capturing stops automatically. The value is limited to 3% of the available system memory at startup (not to exceed 4GB). If a value is not specified, the default packet capture file size is 100MB.
 - Capture last n matching Kilobytes. Enter the number of kilobytes (n) to capture. Any packet received after the memory limit is reached results in the discarding of the oldest saved packet prior to saving the new packet. The saved packets in memory are written to disk when the capture is stopped. The value is limited to 3% of the available system memory at startup (not to exceed 4GB). If a value is not specified, the default packet capture file size is 100MB.
- b. Optional—To truncate the number of bytes saved in each frame, enter a number in the **Save first n bytes of each packet** field. When configured, `pcap` collects, at most, n bytes of packets from each frame when writing to disk. The range is 1 to 65535.
- c. Optional—To specify the number of kilobytes of packets kept in a core image, enter a value in the **Include n K Bytes in core image** field. You can capture packets and include them along with a core image. This is extremely useful if a certain pattern of packets causes the unit to restart unexpectedly. The core image size The value is limited to 3% of the available system memory at startup (not to exceed 4GB). By default, no packets are kept in the core image.
- d. To start the capture, click **Start Capture**. The Start Capture dialog closes. The **Start captures** button in the **Packet Captures** tab is now grayed out because packet capturing is already started.

You do not have to click **Apply** because all changes are applied when you start the packet capture.



4. To stop the capture, click the **Stop capture** button. This button is grayed out if a packet capture is already stopped.
5. To download the capture, click the **Download capture** button. This button is grayed out if no file is available for downloading.

To start, stop, and download packet captures through a browser:

1. Start your Web browser.
2. Enter the URL: `https://appliance_IP_address:8082/PCAP/Statistics` and log in to the appliance as needed. The Packet Capture browser displays.

Packet Capture

Packet capture Statistics

Current state: Stopped
Filtering: Off

Packet capture information:
Packets captured : 181
Bytes captured : 19,549
Packets written : 181
Bytes written : 24,745
Coreimage ram used : 0 B
Packets filtered through : 0

[Start](#) packet capture

[Stop](#) packet capture

[Download](#) packet capture file

3. Select the desired action: **Start packet capture**, **Stop packet capture**, **Download packet capture file**.

You can also use the following URLs to configure these individually:

- ❑ To start packet capturing, use this URL:
`https://ProxySG_IP_address:8082/PCAP/start`
- ❑ To stop packet capturing, use this URL:
`https://ProxySG_IP_address:8082/PCAP/stop`
- ❑ To download packet capturing data, use this URL:
`https://ProxySG_IP_address:8082/PCAP/bluecoat.cap`

Viewing Current Packet Capture Data

Use the following procedures to display current capture information from the ProxySG.

To view current packet capture statistics:

1. Select the **Maintenance > Service Information > Packet Captures** tab.
2. To view the packet capture statistics, click **Show statistics**.

A window opens displaying the statistics on the current packet capture settings. Close the window when you are finished viewing the statistics.

Uploading Packet Capture Data

Use the following command to transfer packet capture data from the ProxySG to an FTP site. You cannot use the Management Console. After uploading is complete, you can analyze the packet capture data.

```
SGOS# pcap transfer ftp://url/path/filename.cap username password
```

Specify a username and password, if the FTP server requires these. The username and password must be recognized by the FTP server.

Core Image Restart Options

This option specifies how much detail is logged to disk when a system is restarted. Although this information is not visible to the user, Blue Coat Technical Support uses it in resolving system problems. The more detail logged, the longer it takes the ProxySG to restart. There are three options:

- ❑ **None**—no system state information is logged. Not recommended.
- ❑ **Context only**—the state of active processes is logged to disk. This is the default.
- ❑ **Full**—A complete dump is logged to disk. Use only when asked to do so by Blue Coat Technical Support.

The default setting of Context only is the optimum balance between restart speed and the information needs of Blue Coat Technical Support in helping to resolve a system problem.

You can also select the number of core images that are retained. The default value is 2; the range is between 1 and 10.

To configure core image restart options:

1. Select **Maintenance > Core Images**.

2. Select a core image restart option.
3. (Optional) Select the number of core images that are retained from the **Number of stored images** drop-down list.
4. Click **Apply**.

Diagnostics: Symantec Customer Experience Program and Monitoring

Every 24 hours, ProxySG appliance transmits a *heartbeat*, which is a periodic message that contains ProxySG statistical data. Besides informing recipients that the device is alive, heartbeats also indicate the health of the appliance. Heartbeats do *not* contain any private information; they only contain aggregate statistics that are invaluable to preemptively diagnose support issues. The daily heartbeat is encrypted and transferred to Blue Coat using HTTPS. You can also have the daily heartbeat messages e-mailed to you by configuring **Event Logging**. The e-mailed content is the same content that is sent to Blue Coat.

You can manage the customer experience program and monitoring settings (heartbeats) from the CLI only as described in the following sections:

To disable heartbeats:

```
SGOS#(config) diagnostics
SGOS#(config diagnostics) heartbeat disable
```

To manually send a heartbeat message:

If you disable automatic heartbeats, you can still manually send a heartbeat message by entering the following commands:

```
SGOS#(config) diagnostics
SGOS#(config diagnostics) send-heartbeat
```

To disable monitoring:

When *monitoring* is enabled (it is enabled by default), Blue Coat receives encrypted information over HTTPS whenever the appliance is rebooted. Like the heartbeat, the data sent does *not* contain any private information; it contains restart summaries and daily heartbeats. This allows the tracking of ProxySG unexpected restarts because of system issues, and allows Blue Coat to address system issues preemptively. To disable monitoring, enter the following commands:

```
SGOS#(config) diagnostics  
SGOS#(config diagnostics) monitor disable
```

To enable heartbeats and/or monitoring:

If you have disabled heartbeats and/or monitoring, you can re-enable them by entering the following commands:

```
SGOS#(config diagnostics) heartbeat enable  
SGOS#(config diagnostics) monitor enable
```


Diagnostic Reporting (CPU Monitoring)

You can enable CPU monitoring whenever you want to see the percentage of CPU being used by specific functional groups. For example, if you look at the CPU consumption and notice that compression/decompression is consuming most of the CPU, you can change your policy to compress/decompress more selectively.


Note: CPU monitoring uses about 2-3% CPU when enabled, and so is disabled by default.

To configure and view CPU monitoring:

1. Select **Statistics > Advanced**.



2. Click the **Diagnostics** link. A list of links to Diagnostic URLs displays.



3. To enable CPU monitoring, click the **Start the CPU Monitor** link; to disable it, click the **Stop the CPU Monitor** link.
4. To view CPU monitoring statistics, click the CPU Monitor statistics link. You can also click this link from either of the windows described in Step 3.

Configure Auto Refresh Interval for Monitoring Statistics

You can configure the interval at which CPU monitoring statistics refresh in the browser. Enter the CLI command:

```
#(config diagnostics) cpu-monitor interval seconds
```

Notes

- ❑ The total percentages displayed on the CPU Monitor Statistics page do not always add up because the display only shows those functional groups that are using 1% or more of the CPU processing cycles.
- ❑ The `SGOS#(config) show cpu` and `SGOS#(config diagnostics) view cpu-monitor` commands might sometimes display CPU statistics that differ by about 2-3%. This occurs because different measurement techniques are used for the two displays.

Chapter 76: XML Protocol

The XML realm uses a SOAP 1.2 based protocol for the Blue Coat supported protocol.

This section includes the following topics:

- ❑ [Section A: "Authenticate Request" on page 1430](#)
- ❑ [Section B: "Authenticate Response" on page 1432](#)
- ❑ [Section C: "Authorize Request" on page 1434](#)
- ❑ [Section D: "Authorize Response" on page 1435](#)

Note: Examples in this chapter refer to an XML schema. Refer to the *SGOS 6.5 Release Notes* for the location of this file.

Section A: Authenticate Request

GET Method (User Credentials in Request)

If the user credentials are not set in the HTTP headers, the username and password are added to the query. The name of the username parameter is configured in the realm. The groups and attributes of interest are only included if the realm is configured to include them.

```
http://<server hostname>:<server port>/<authenticate service path>?<username parameter name>=<username>&password=<password>[&group=<group 1>&group=<group 2>...&attribute=<attribute 1>&attribute=<attribute 2>]
```

GET Method (User Credentials in Headers)

If the user credentials are in the HTTP headers, the password is not added to the query.

```
http://<server hostname>:<server port>/<authenticate service path>/authenticate?<username parameter name>=<username>[&group=<group 1>&group=<group 2>...&attribute=<attribute 1>&attribute=<attribute 2>]
```

POST Method (User Credentials in Request)

The parameter name of the username is configured in the realm. The groups and attributes of interest are included only if the realm is configured to include them.

```
<?xml version='1.0'encoding="UTF-8" ?>
<env:Envelope
xmlns:env="http://www.w3.org/2003/05/soap-envelope">
  <env:Body env:encodingStyle="http://www.w3.org/2003/05/soap-encoding"
    xmlns:enc="http://www.w3.org/2003/05/soap-encoding">
    <m:authenticate
      xmlns:m="http://www.bluecoat.com/xmlns/xml-realm/1.0">
      <m:username>Username</m:username>
      <m:password>password</m:password>
      <m:groups enc:arraySize="*" enc:itemType="xsd:string">
        <m:group>group1</m:group>
        <m:group>group2</m:group>
      </m:groups>
      <m:attributes enc:arraySize="*" enc:itemType="xsd:string">
        <m:attribute>attribute1</m:attribute>
        <m:attribute>attribute2</m:attribute>
      </m:attributes>
    </m:authenticate>
  </env:Body>
</env:Envelope>
```

POST Method (User Credentials in Headers)

If the user credentials are in the HTTP headers, the password is not added to the request.

```
<?xml version='1.0' encoding="UTF-8" ?>
<env:Envelope
xmlns:env="http://www.w3.org/2003/05/soap-envelope">
  <env:Body
env:encodingStyle="http://www.w3.org/2003/05/soap-encoding">
    <m:authenticate
xmlns:m="http://www.bluecoat.com/xmlns/xml-realm/1.0">
      <m:username>Username</m:username>
      <m:challenge-state>challenge state</m:challenge-state>
      <m:groups enc:arraySize="*" enc:itemType="xsd:string">
        <m:group>group1</m:group>
        <m:group>group2</m:group>
      </m:groups>
      <m:attributes enc:arraySize="*" enc:itemType="xsd:string">
        <m:attribute>attributel</m:attribute>
        <m:attribute>attribute2</m:attribute>
      </m:attributes>
    </m:authenticate>
  </env:Body>
</env:Envelope>
```

Section B: Authenticate Response

Success

All of the response fields except `full-username` are optional. The intersection of the groups of interest and the groups that the user is in are returned in the `groups` element. The attributes of interest for the user are returned in a flattened two dimensional array of attribute names and values.

```
<?xml version='1.0' encoding="UTF-8" ?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope">
  <env:Body
    env:encodingStyle="http://www.w3.org/2003/05/soap-encoding">
    <m:authenticate-response
      xmlns:m="http://www.bluecoat.com/xmlns/xml-realm/1.0">
      <m:full-username>full-username</m:full-username>
      <m:groups enc:arraySize="*" enc:itemType="xsd:string">
        <m:group>group2</m:group>
      </m:groups>
      <m:attribute-values enc:arraySize="* 2"
        enc:itemType="xsd:string">
        <m:item>attribute2</m:item>
        <m:item>value2a</m:item>
        <m:item>attribute2</m:item>
        <m:item>value2b</m:item>
        <m:item>attribute2</m:item>
        <m:item>value2c</m:item>
      </m:attribute-values>
    </m:authenticate-response>
  </env:Body>
</env:Envelope>
```

Failed/Denied

The failed response includes a text description of the failure that becomes the text description of the error reported to the user. The fault-code is one of a set of SGOS authentication errors that can be returned from the responder. The codes are returned as strings, but are part of an enumeration declared in the schema for the protocol. Only codes in this list are acceptable.

```
account_disabled
account_restricted
credentials_mismatch
general_authentication_error
expired_credentials
account_locked_out
account_must_change_password
offbox_server_down
general_authorization_error
unknown_error
```



```
<?xml version='1.0' encoding="UTF-8" ?>
<env:Envelope
xmlns:env="http://www.w3.org/2003/05/soap-envelope">
  <env:Body>
    <env:Fault>
      <env:Code>
        <env:Value>env:Sender</env:Value>
      </env:Code>
      <env:Reason>
        <env:Text xml:lang="en-US">Bad username or password</env:Text>
      </env:Reason>
      <env:Detail>
        <e:realm-fault
xmlns:e="http://www.bluecoat.com/xmlns/xml-realm/1.0">
          <e:fault-code>general_authentication_error</e:fault-code>
        <e:realm-fault>
      </env:Detail>
    </env:Fault>
  </env:Body>
</env:Envelope>
```

Section C: Authorize Request

The groups and attributes of interest for the user are embedded in the request if they are configured to be included. The XML responder must not require credentials for authorization requests.

GET Method

```
http://<server hostname>:<server port>/<authorize service
path>?<username parameter
name>=<username>[&group=<group1>&group=<group2>...&attribute=<attribute1
>&...]
```

POST Method

```
<?xml version='1.0' encoding="UTF-8" ?>
<env:Envelope
xmlns:env="http://www.w3.org/2003/05/soap-envelope">
  <env:Body
env:encodingStyle="http://www.w3.org/2003/05/soap-encoding"
xmlns:enc="http://www.w3.org/2003/05/soap-encoding">
    <m:authorize
xmlns:m="http://www.bluecoat.com/soap/xmlns/xml-realm/1.0">
      <m:username>Username</m:username>
      <m:groups enc:arraySize="*" enc:itemType="xsd:string">
        <m:group>group1</m:group>
        <m:group>group2</m:group>
      </m:groups>
      <m:attributes enc:arraySize="*" enc:itemType="xsd:string">
        <m:attribute>attribute1</m:attribute>
        <m:attribute>attribute2</m:attribute>
      </m:attributes>
    </m:authorize>
  </env:Body>
</env:Envelope>
```

Section D: Authorize Response

Success

Only applicable groups and attributes are returned. Multi-valued attributes are returned by multiple instances of the same attribute name.

```
<?xml version='1.0' encoding="UTF-8" ?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope">
  <env:Body
    env:encodingStyle="http://www.w3.org/2003/05/soap-encoding"
    xmlns:enc="http://www.w3.org/2003/05/soap-encoding">
    <m:authorize-response
      xmlns:m="http://www.bluecoat.com/xmlns/xml-realm/1.0">
      <m:groups enc:arraySize="*" enc:itemType="xsd:string">
        <m:group>group2</m:group>
      </m:groups>
      <m:attribute-values enc:arraySize="* 2"
        enc:itemType="xsd:string">
        <m:item>attribute2</m:item>
        <m:item>value2a</m:item>
        <m:item>attribute2</m:item>
        <m:item>value2b</m:item>
        <m:item>attribute2</m:item>
        <m:item>value2c</m:item>
      </m:attribute-values>
    </m:authorize-response>
  </env:Body>
</env:Envelope>
```

Failed

```
<?xml version='1.0'encoding="UTF-8" ?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope">
  <env:Body>
    <env:Fault>
      <env:Code>
        <env:Value>env:Receiver</env:Value>
      </env:Code>
      <env:Reason>
        <env:Text xml:lang="en-US">Could not contact LDAP server</
env:Text>
      </env:Reason>
      <env:Detail>
        <e:realm-fault
          xmlns:e="http://www.bluecoat.com/xmlns/xml-realm/1.0">
          <e:fault-code>offbox_server_down</e:fault-code>
        </e:realm-fault>
      </env:Detail>
    </env:Fault>
  </env:Body>
</env:Envelope>
```

